COMPUTATIONAL ISSUES
IN
REGULARIZED RESTORATION
USING
LARGE IMAGING OPERATORS

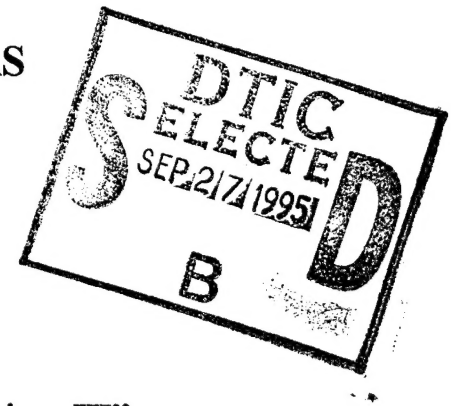

Based on a presentation at:
SPIE Technical Conference "Inverse Optics III"
Orlando, Florida, 4-5 April 1994

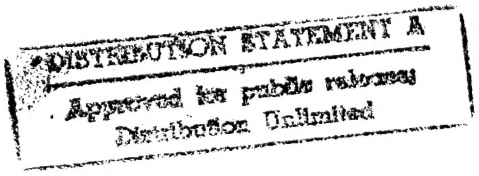

John B. Abbiss
Titan Spectron, 1582 Parkway Loop, Suite B., Tustin, California 92680-6505

Jeffery C. Allen
Harper J. Whitehouse
Naval Command, Control and Ocean Surveillance Center
RDT & E Division, San Diego, California 92152-5000

ACKNOWLEDGEMENTS

This research was supported by the Innovative Science and Technology Office of the Ballistic Missile Defense Organization, through the Office of Naval Research. The program was managed by Dr. Keith Bromley.

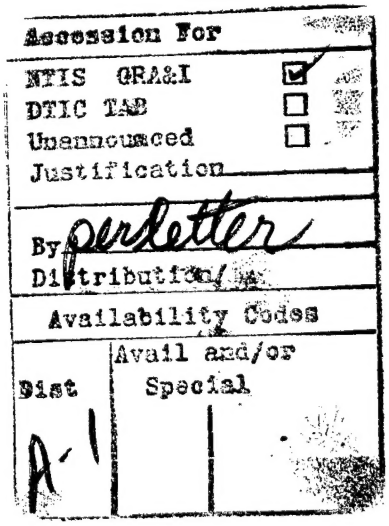

SUMMARY

We consider regularized linear algebraic methods for the restoration of large images (perhaps 1000 x 1000 pixels) derived from space-based surveillance systems. The associated computational burdens are assessed. Methods which accelerate the restoration process by exploiting structure within the imaging matrix are analyzed. It is shown that, if relaxation of the support of the reconstruction to that of the original image can be tolerated, a technique based on expansion of the imaging matrix to circulant form can be applied which is significantly faster and makes much smaller demands on storage resources.

1. The Discretized Problem

We wish to form an estimate of an object f, given the image g and imaging operator A. We assume additive noise is present and write, symbolically,

$$g = Af + n \tag{1}$$

Typically, A is a strongly smoothing operator - for example, a convolution in the case of a space-invariant point-spread function - and noise introduces a destabilizing effect on any attempt at a solution. In the continuous form of the problem, equation (1) becomes a Fredholm equation of the first kind, and a singular function analysis can be used[1] to identify the precise nature of this difficulty. In the discrete case, A is a matrix, f, g and n are vectors, and the singular value decomposition (SVD) of A can be used in a similar way[1]. We shall see that the SVD also provides an elegant technique for deriving a solution.

In the presence of noise, an exact solution to equation (1) is impossible to achieve. As a means of choosing an approximation from the infinite set $\left\{\hat{f}\right\}$ of estimates, we might require, for example, that the Euclidean distance between the image of the estimate and the given image data be a minimum.

Thus we would minimize $\left\|A\hat{f} - g\right\|_2^2$ and obtain the normal equations

$$A^H A f_{est} = A^H g \quad ,$$

where A^H is the Hermitian conjugate of A; inversion of $A^H A$ (assuming that it has an inverse) then yields the estimate f_{est}. It can be shown, however, that this estimator is still extremely sensitive to noise, small perturbations in the data causing large changes in the solution. The problem is intrinsically ill-posed. It can be converted to a well-posed one by the methods of regularization theory[2]. We add an 'energy' constraint,

$$\left\| \hat{f} \right\|_2^2 \leq E^2 \quad ,$$

combine it with the previous constraint, and minimize

$$\left\| A\hat{f} - g \right\|_2^2 + \beta \left\| \hat{f} \right\|_2^2 \quad ,$$

where β, the regularization parameter, plays the role of a Lagrange multiplier. The solution becomes

$$f_\beta = \left(A^H A + \beta I \right)^{-1} A^H g$$

or

$$f_\beta = A_\beta^+ \, g \quad ,$$

where $A_\beta^+ = \left(A^H A + \beta I \right)^{-1} A^H$ is the regularized pseudoinverse of A, by analogy with the Moore-Penrose pseudoinverse[3] A^+. The optimum value for β will depend on the noise level and can be determined in one of several ways[4,5].

A_β^+ could be computed directly by forming $A^H A$ and carrying out the required matrix inversion. This procedure loses accuracy, however, and SVD is the method of choice[6]. Let A be an $n \times n$ matrix. (The formalism is easily extended to the general $m \times n$ case.) We have

$$A = U \Sigma V^H \quad ,$$

where $$U^H U = V^H V = V V^H = I_n$$

and $$\Sigma = diag\ (\sigma_1, \ldots, \sigma_n) \ .$$

The singular values $\{\sigma_i\}$ are assumed to have been arranged in descending order of magnitude:

$$\sigma_1 \geq \sigma_2 \geq .... \geq \sigma_n \geq 0 \ .$$

Then we find $$f_\beta = V\Sigma_\beta^+ U^H g \ ,$$

where $$\Sigma_\beta^+ = diag\left(..., \frac{\sigma_i}{\sigma_i^2 + \beta}, ... \right) \ .$$

SVD is, however, computationally expensive, involving $O(10N^3)$ operations for an NxN matrix. A Matlab experiment on a 66 MHz 80486 PC produced the following results:

A-Matrix	Time to compute SVD	Output file	Number of operations
512 x 512	11m 16s	6.3 Mb	1.33 Gflops

The *A*-matrix associated with a 23x23 image would be of approximately this size.

2. Reducing the Computational Burden

In certain cases, the computational effort can be reduced significantly. For example, if the point-spread function is stationary across the field of view and unchanging in time, and if the image size does not vary, the SVD can be precomputed and stored.

The remaining (on-line) computation is

$$f_\beta = A_\beta^+ g = V\Sigma_\beta^+ U^H g$$

an $O(N^2)$ procedure, where N is the number of pixels in g.

The matrix-vector multiplication effort can be further reduced by exploiting structure in A (and hence in A_β^+).

For stationary point-spread function and equal sampling intervals in image g and reconstruction f_β, A is Toeplitz (i.e., $A = \left[a_{j-i}\right]$) in the one-dimensional case. In two dimensions, image and reconstruction must be reconfigured from matrices into vectors in some way; for example, by column-wise mapping. A then becomes block Toeplitz with Toeplitz blocks, consisting of Toeplitz submatrices arranged within A in Toeplitz fashion.

Figure 1 is a Matlab meshplot of an 18 x 18 test object consisting entirely of 1's or 0's located at the intersections of the grid lines. The 8 x 8 truncated Gaussian point-spread function is shown in Figure 2. The blurred image of Figure 3 obtained with this point-spread function then occupies 25 x 25 pixels; note that no noise has been added to this image. The 625 x 324 imaging matrix A, which relates the column-mapped object and image vectors, is shown in Figure 4 in perspective and in contour form in Figure 5. Its Toeplitz symmetries are clearly discernible.

Figure 1

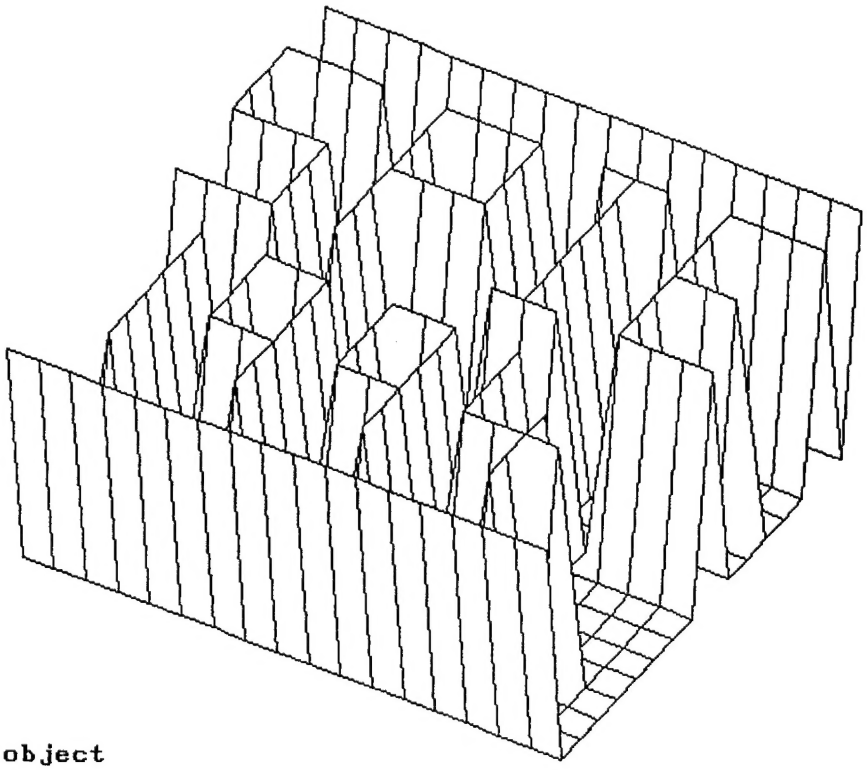

18x18 object

Figure 2

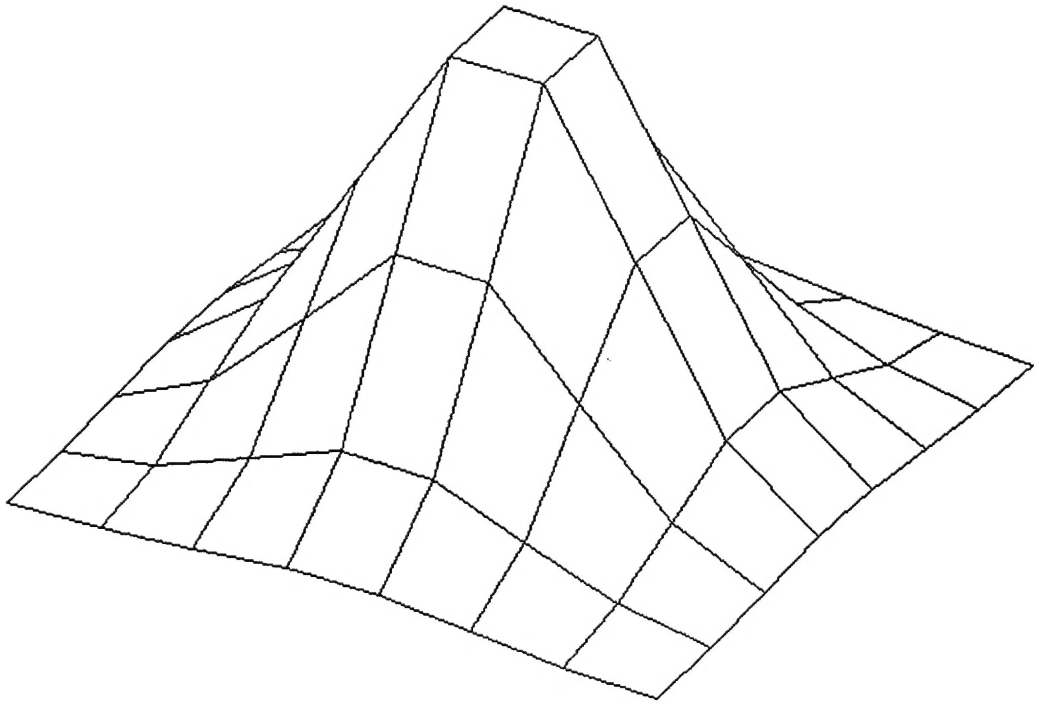

8x8 Gaussian point-spread function

Figure 3

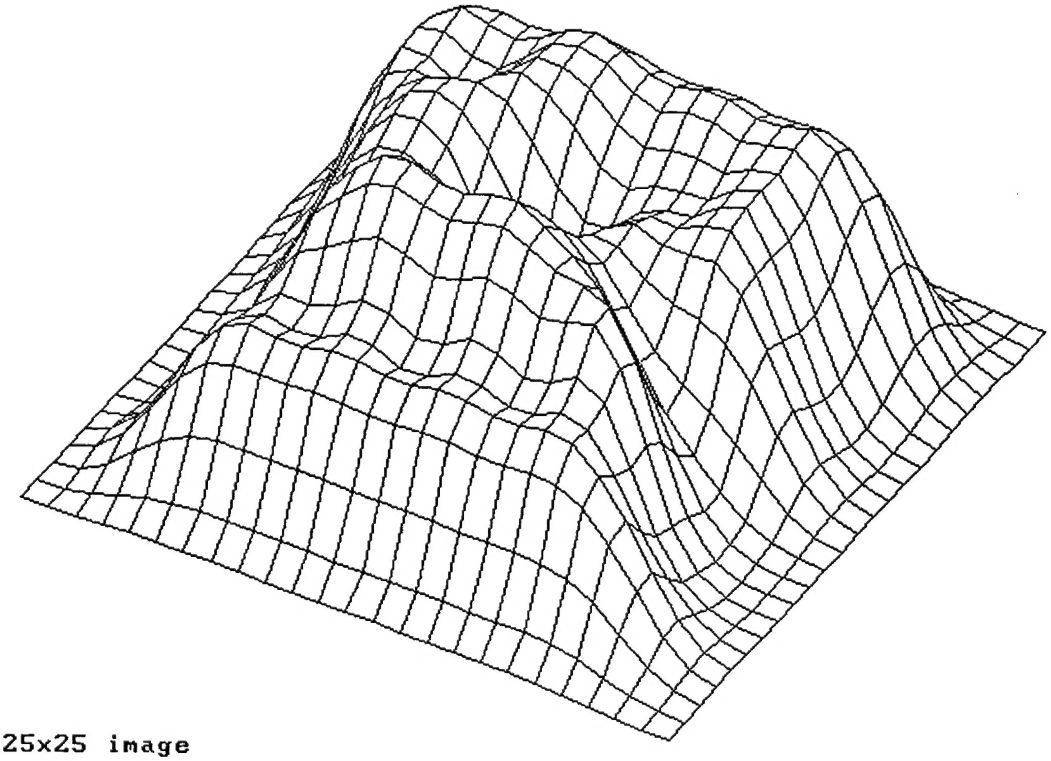

25x25 image

Figure 4

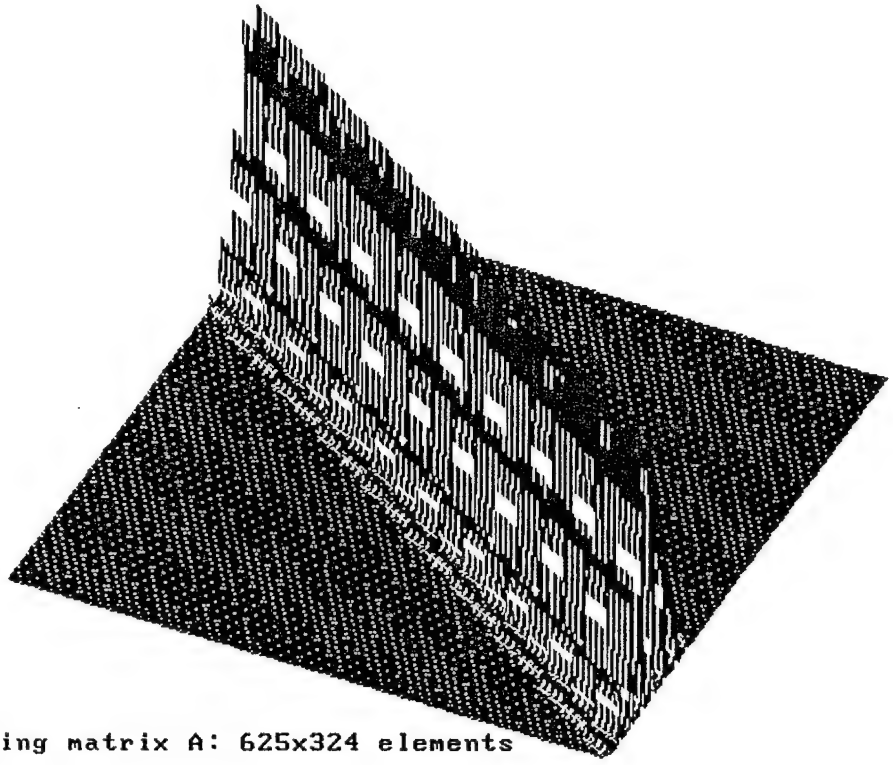

Figure 5

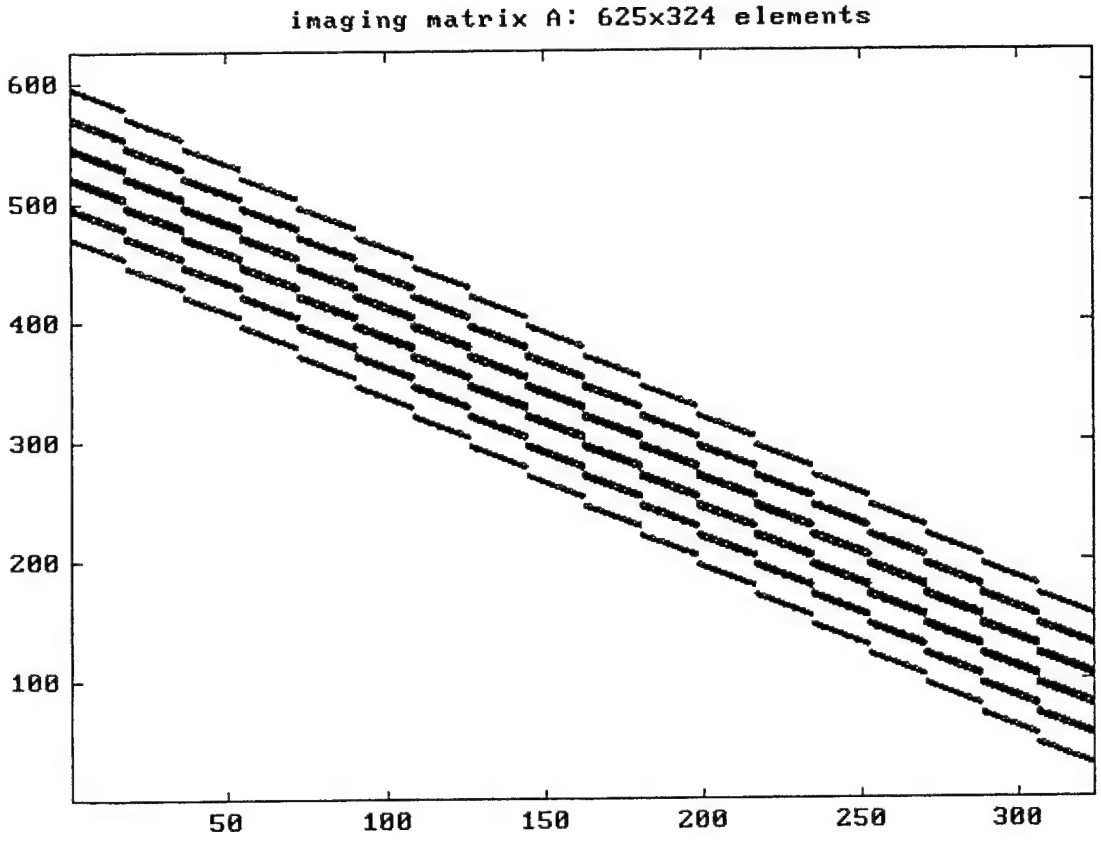
imaging matrix A: 625x324 elements
600
500
400
300
200
100
50
100
150
200
250
300

Figure 6

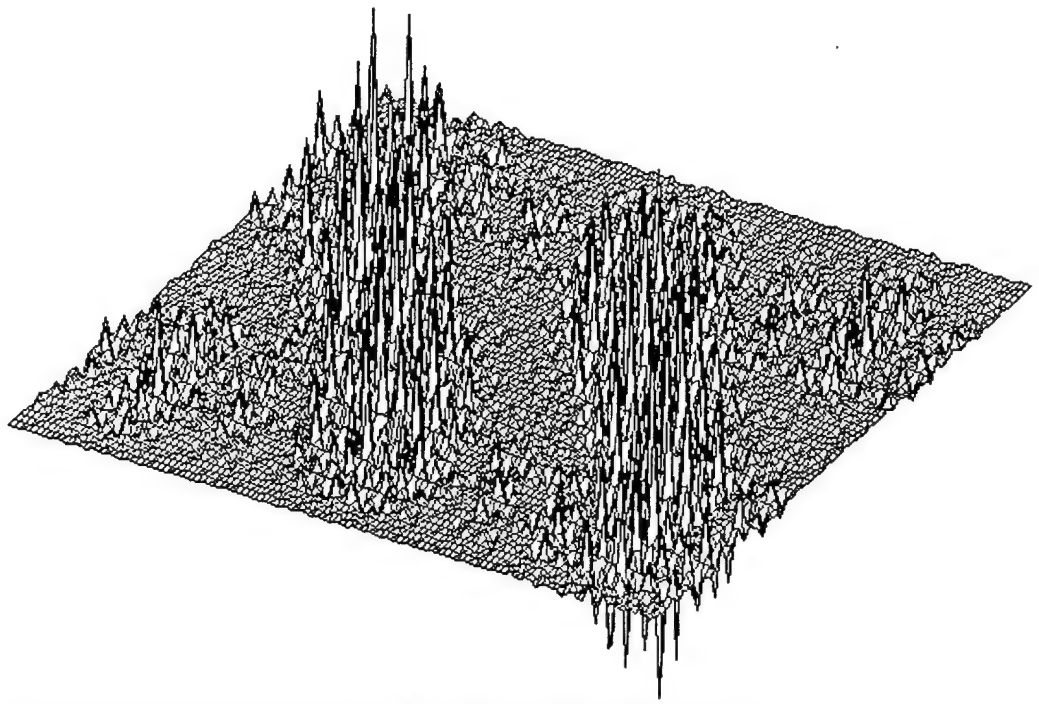

Figure 7

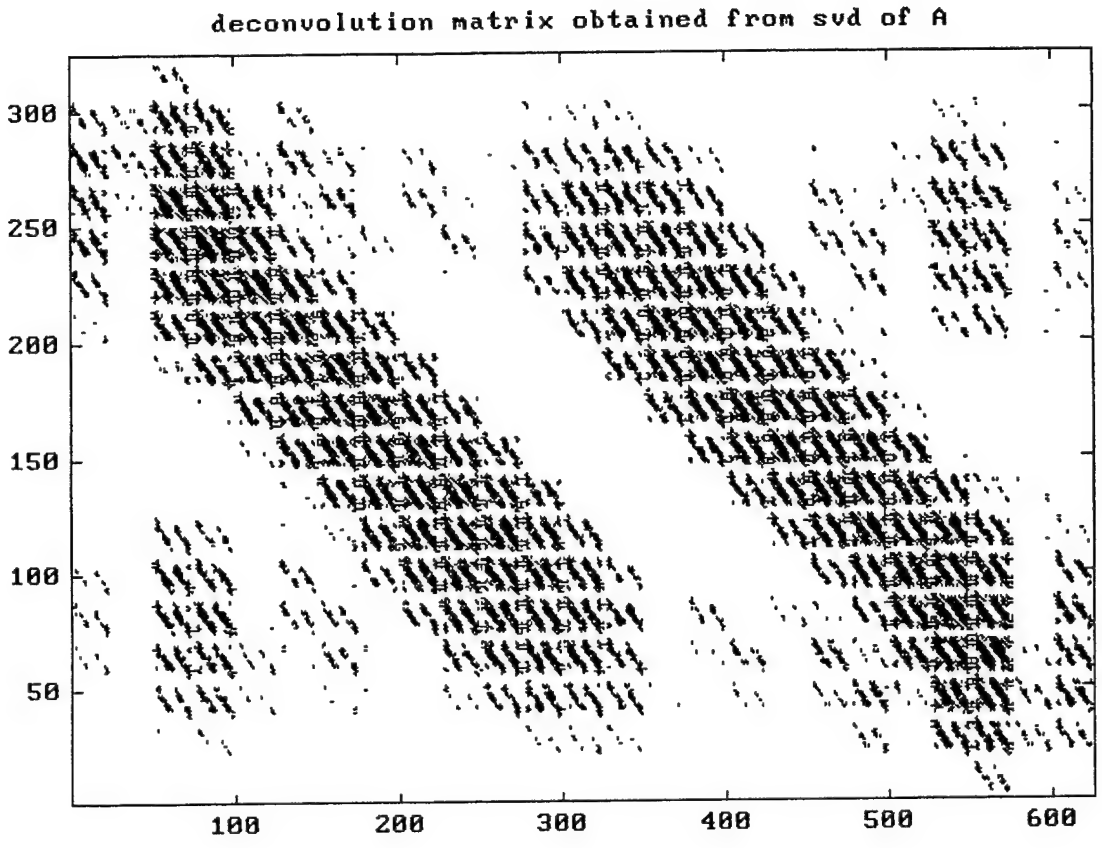
deconvolution matrix obtained from svd of A
300
250
200
150
100
50
100
200
300
400
500
600

Figure 8

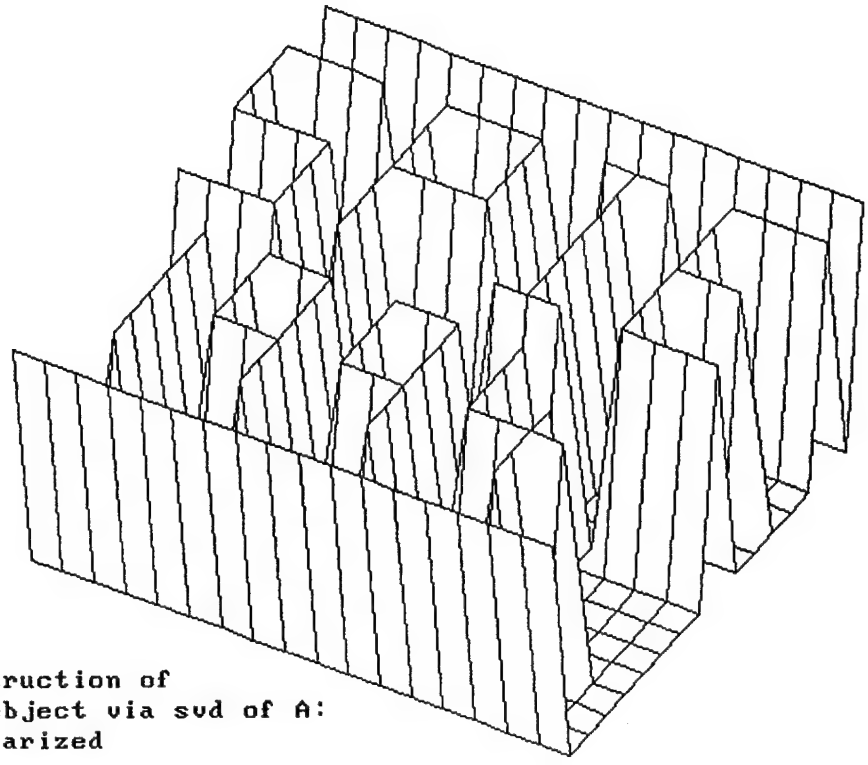

3. Exploiting Symmetries

Symmetries in A_β^+, deriving from the Toeplitz properties of A, are clearly evident in Figure 6 and 7. Such Toeplitz structure can be exploited using the concept of displacement rank[7]. The displacement rank of a general matrix B is the smallest number α for which we can write

$$\mathrm{B} = \sum_{k=1}^{\alpha} L(x_k)\, U(y_k) \quad ,$$

where L and U are lower and upper and triangular Toeplitz matrices. The x_k and y_k can be obtained by SVD of the 'displaced difference' B-ZBZ$^{\mathrm{T}}$ of the matrix, where Z is the downshift matrix

$$\begin{bmatrix} 0 & 0 & 0 & \ldots & 0 & 0 \\ 1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & 1 & 0 & \ldots & 0 & 0 \\ & & & \ldots & & \\ 0 & 0 & 0 & \ldots & 1 & 0 \end{bmatrix} \quad .$$

For A_β^+, we expand the SVD of the displaced difference in the following form[8]:

$$A_\beta^+ - ZA_\beta^+ Z^T = \sum_{k=1}^{\alpha} s_k p_k q_k^H ,$$

where the p_k and q_k are the left and right orthonormal singular vectors and the singular values s_k are indexed in decreasing order. This expression leads to the following displacement-rank expansion for A_β^+:

$$A_\beta^+ = \sum_{k=1}^{\alpha} s_k L(p_k) U(q_k) \quad .$$

Here $L(p_k)$ denotes a lower Toeplitz matrix with first column p_k , and (q_k) an upper Toeplitz matrix with first row q_k^H.

This expansion enables the matrix-vector product $A_\beta^+ g$ to be performed by FFT methods, thus conferring a potential gain in speed and storage requirements. However, an expensive SVD is required to find the p_k and q_k and in practice α may not be an especially small number.

We can achieve much greater computational efficiency, and also develop a procedure with wide applicability, by keeping the pixel sizes in the image and reconstruction spaces equal and by allowing the support of the reconstruction to expand to that of the image.

A can then be made a circulant[9], in which the rows and columns are cyclical, and in which any one column (or row) contains all the information needed to construct the complete matrix A. The elements of the circulant of order n are related by the equation $A = \left[a_{j-i+1}\right]$, where the subscript is taken modulo n. Its key property for the present purpose is that the Fourier transform diagonalizes a circulant. We can write

$$A = F^H \Lambda F \quad ,$$

where $$F^H = \frac{1}{\sqrt{n}} \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & w & w^2 & \dots & w^{n-1} \\ 1 & w^2 & w^4 & \dots & w^{2(n-1)} \\ & & & \dots & \\ 1 & w^{n-1} & w^{2(n-1)} & \dots & w^{(n-1)^2} \end{bmatrix}, \quad w = \exp(i2\pi / n)$$

and $\Lambda = diag(\lambda_1,....,\lambda_n)$.

Then from $$f_\beta = \left(A^H A + \beta I\right)^{-1} A^H g$$

we find $$f_\beta = F^H \Lambda_\beta^+ F g \quad ,$$

where $$\Lambda_\beta^+ = diag\left(...,\frac{\lambda_i^*}{|\lambda_i|^2 + \beta},...\right) \quad .$$

We now need only 1-D FFTs. The λ_i are given by the FFT of the first column of A and the reconstructed image is computed as

$$f_\beta = \frac{1}{\sqrt{n}} ifft\left(\Lambda_\beta^+ fft(g)\right) \quad ,$$

which is an $O(3n)\log_2 n$ procedure.

The penalty is that the support of f_β has been relaxed - noise can leak into 'zero' areas. For images much larger than the effective point-spread function, however, this is not a serious problem.

Figures 9 and 10 show the circulant form of the imaging matrix for the same point-spread function and image size as in the earlier example, and should be compared with Figures 4 and 5. The reconstruction now expands to 25 x 25 pixels. It can easily be shown that the deconvolution matrix A_β^+, displayed in Figures 11 and 12, is also a circulant, although much less sparse. In the noiseless case, the reconstruction shown in Figure 13 is obtained with $\beta = 0$. When Gaussian noise is added to the image with a standard deviation of 10^{-6} relative to the local pixel value, the reconstruction is as shown in Figure 14. At a noise

level of 10^{-5}, the object is no longer recognizable; see Figure 15. By setting the regularization parameter equal to 10^{-10}, however, the image is acceptably restored, as can be seen in Figure 16. Note also that the leakage of energy outside the true object support is now very limited.

Figure 9

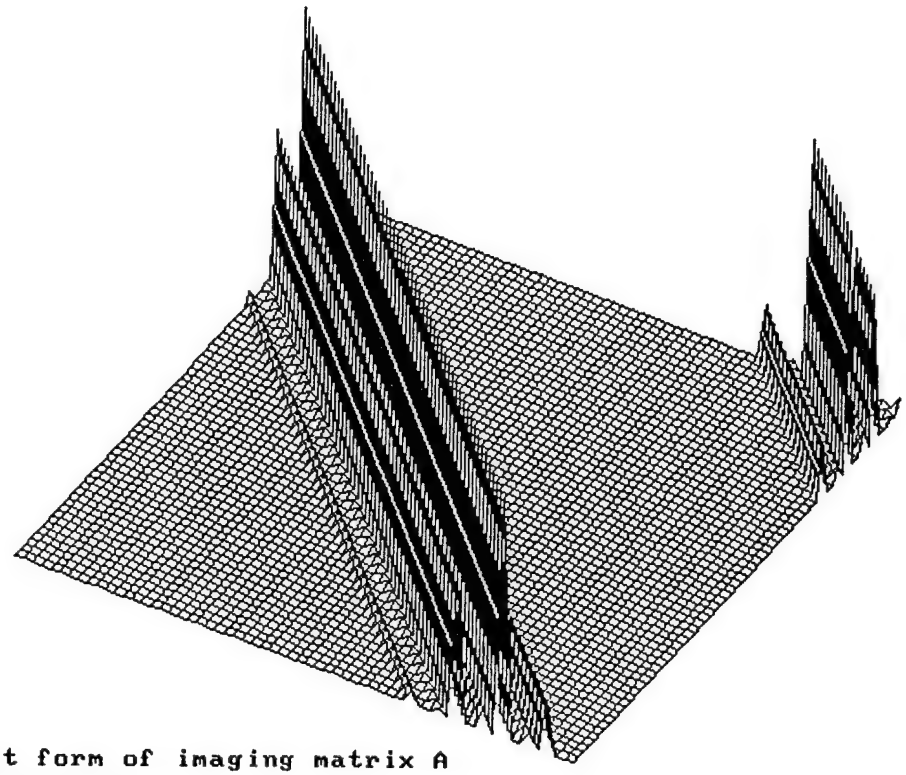

circulant form of imaging matrix A
625x625 elements

Figure 10

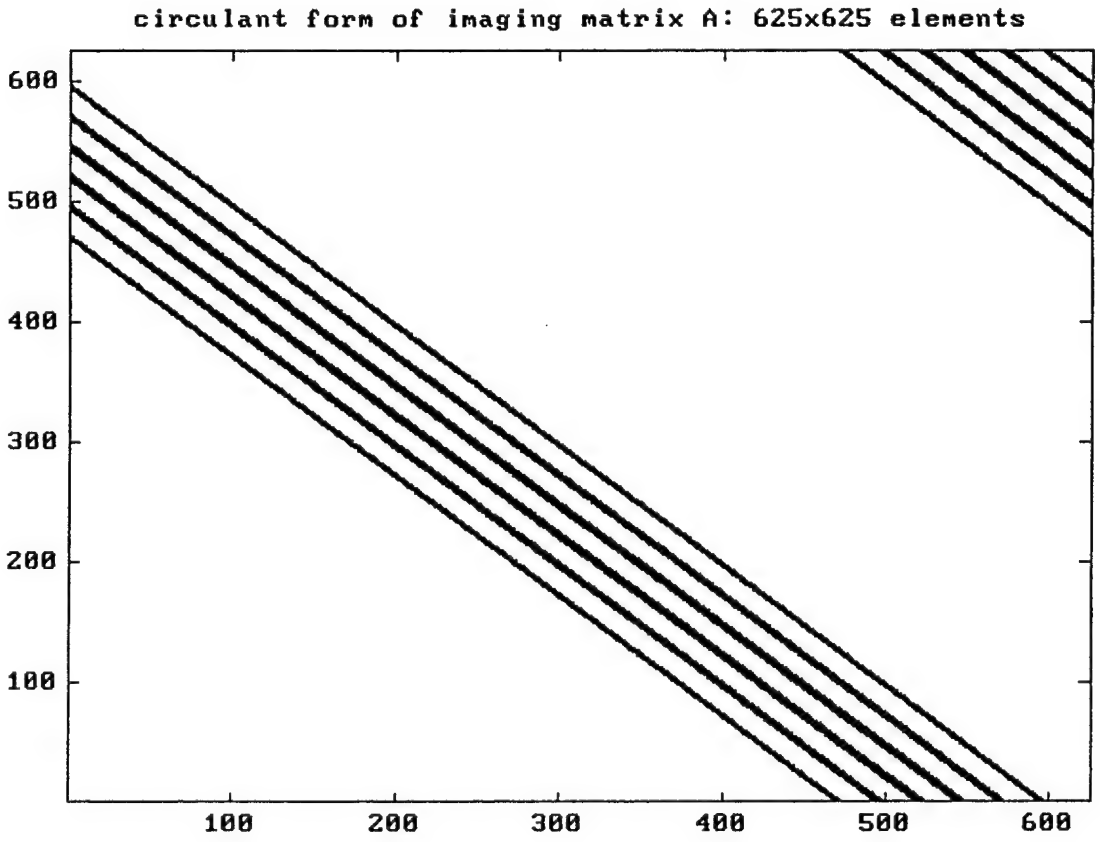
circulant form of imaging matrix A: 625x625 elements
600
500
400
300
200
100
100
200
300
400
500
600

Figure 11

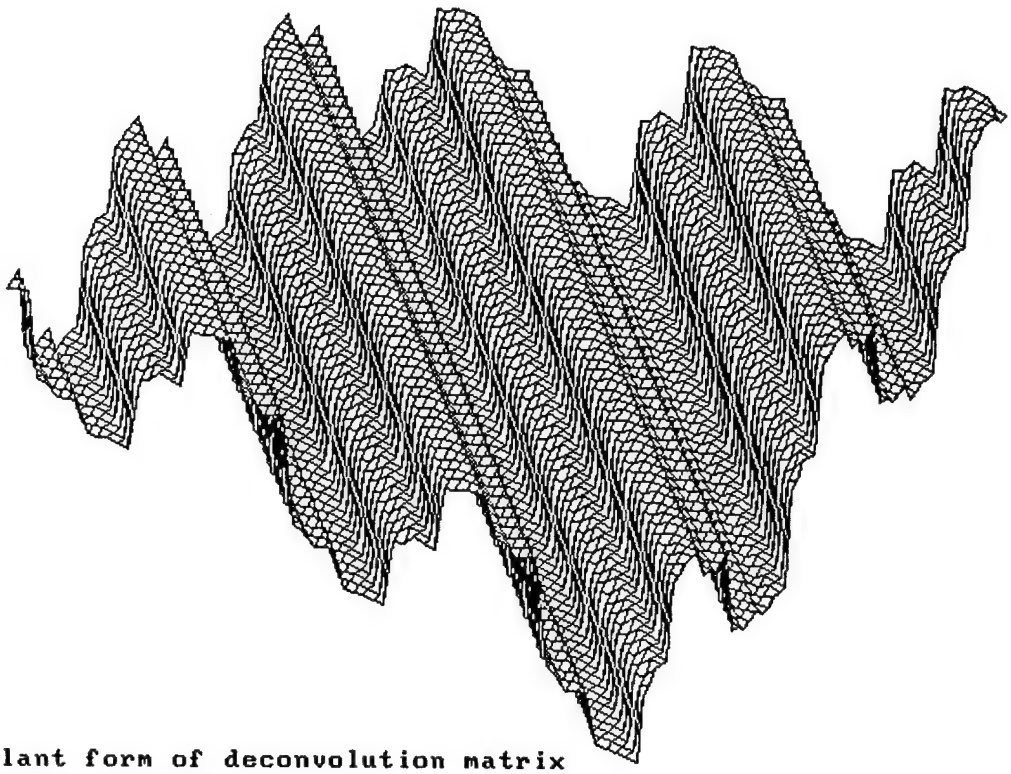

circulant form of deconvolution matrix
625x625 elements

Figure 12

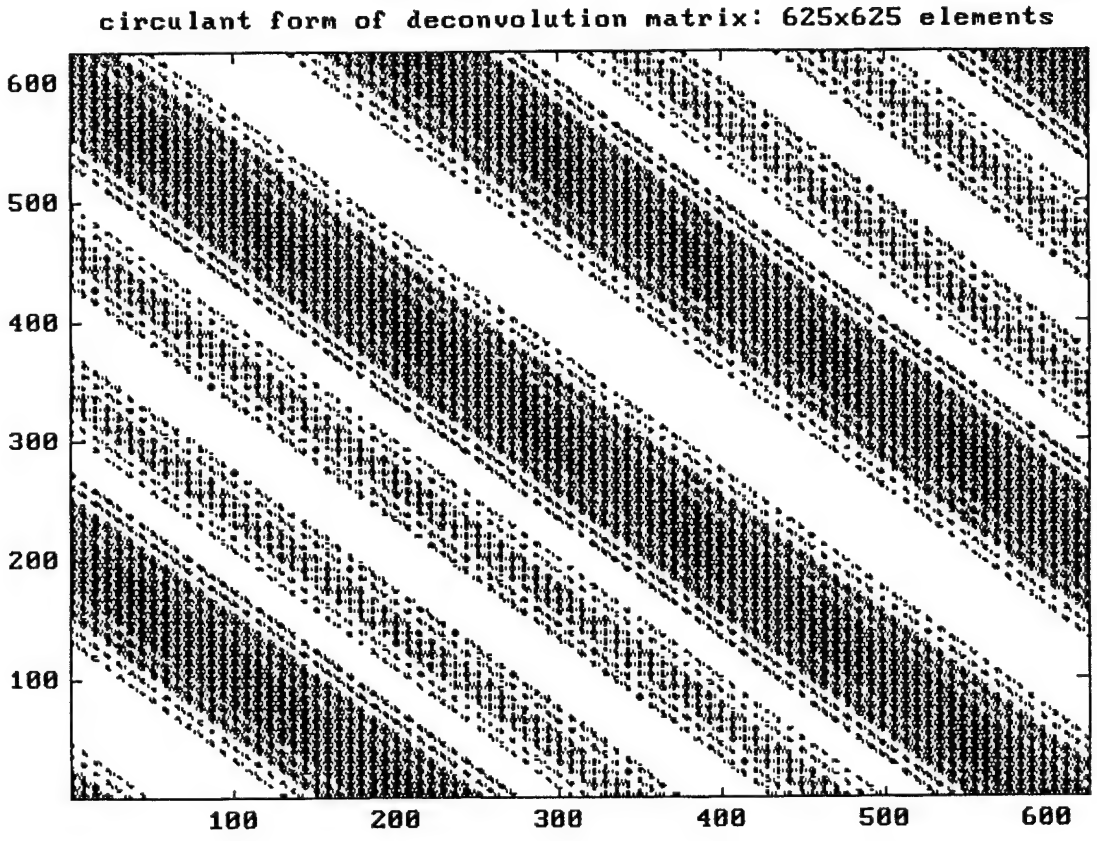
circulant form of deconvolution matrix: 625x625 elements
600
500
400
300
200
100
100
200
300
400
500
600

Figure 13

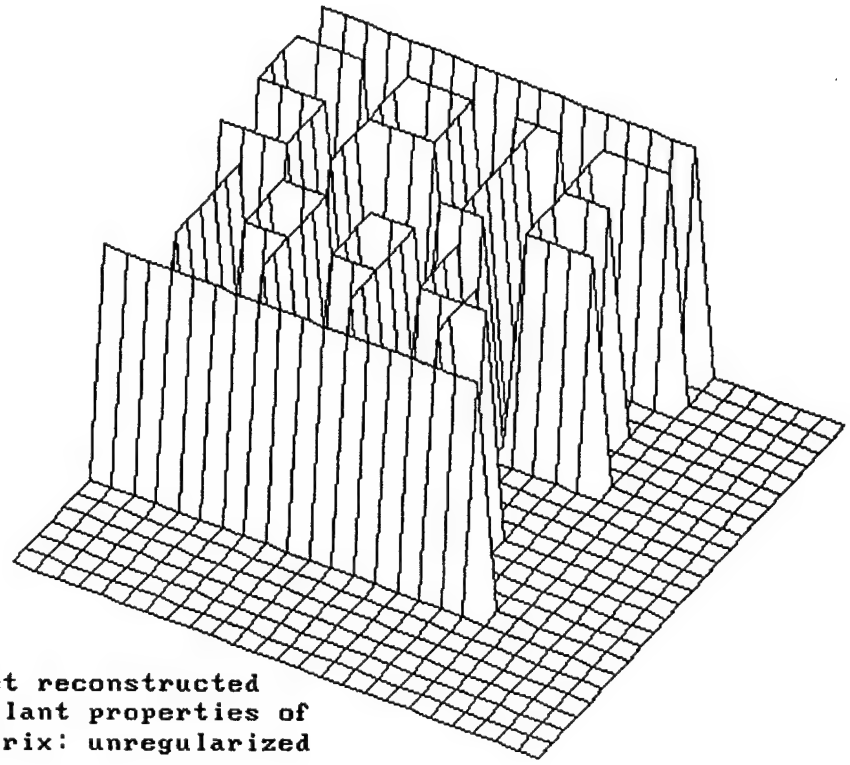

18x18 object reconstructed
using circulant properties of
imaging matrix: unregularized

Figure 14

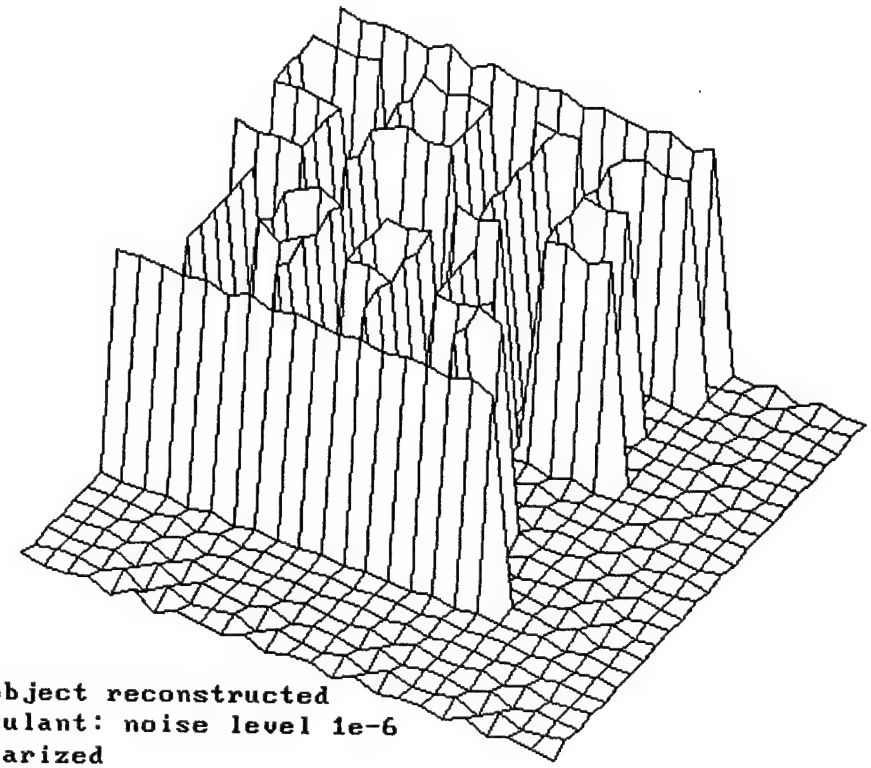

18x18 object reconstructed
by circulant: noise level 1e-6
unregularized

Figure 15

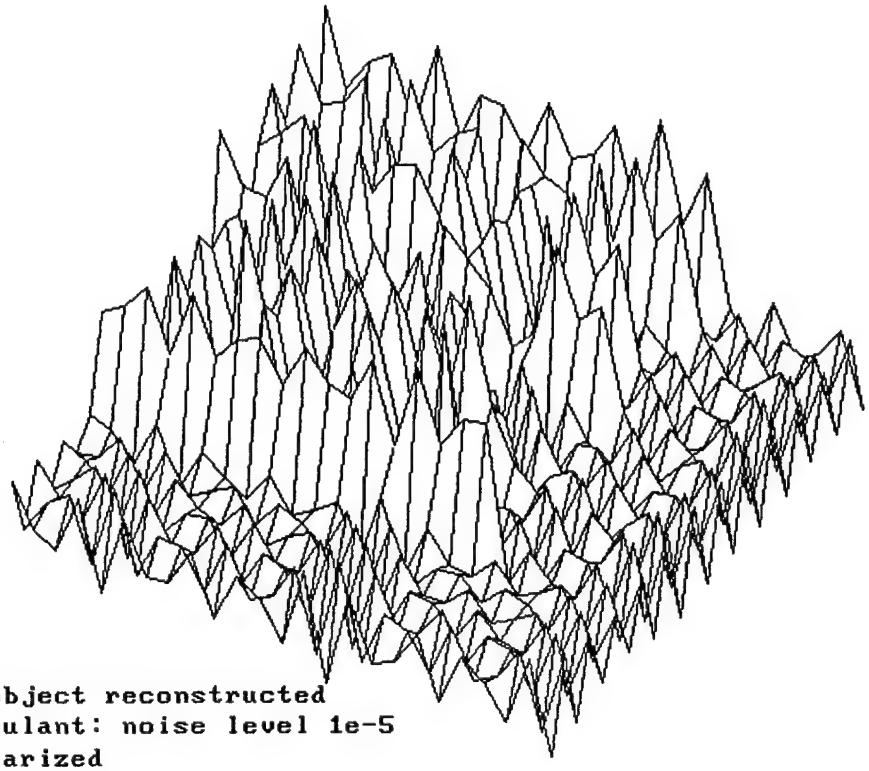

18x18 object reconstructed
by circulant: noise level 1e-5
unregularized

Figure 16

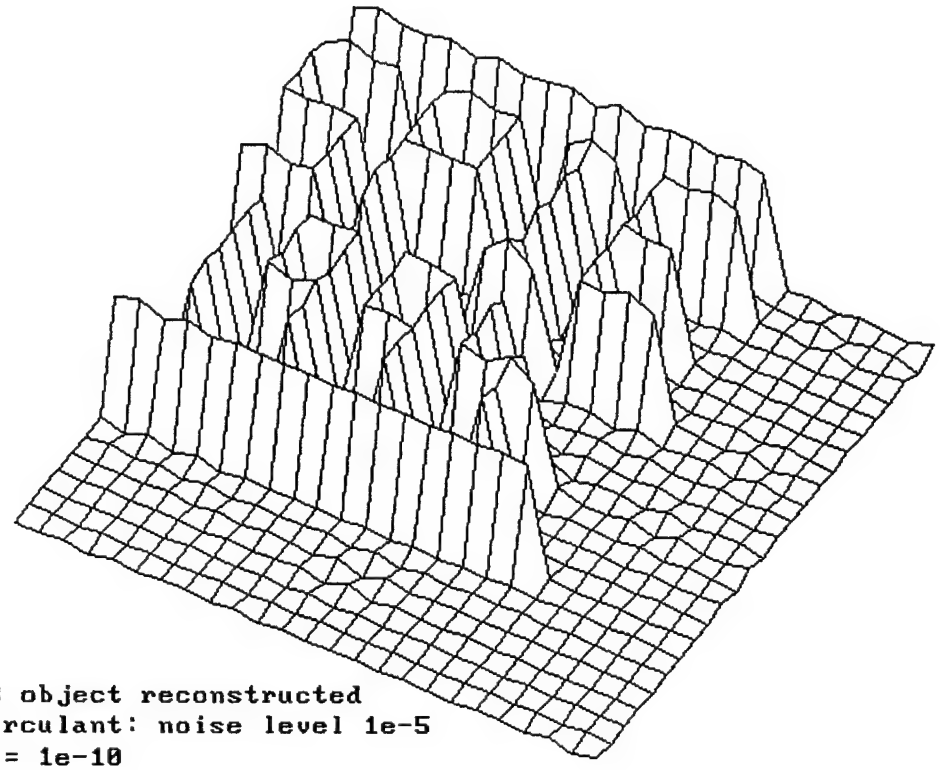

18x18 object reconstructed
by circulant: noise level 1e-5
beta = 1e-10

For a space-variant point-spread function, direct expansion of A to circulant form is no longer possible. However, if variation across the image is small, a displacement-rank expansion of A_β^+ into sums of products of upper triangular Toeplitz and circulant matrices can be made[10]:

$$A_\beta^+ = C_o + \sum_{k=1}^{\alpha} L(x_k)\, C^H(y_k) \quad .$$

The x_k and y_k are now obtained by SVD of the cyclic displacement of A_β^+:

$$A_\beta^+ - E A_\beta^+ E^T = \sum_{k=1}^{\alpha} x_k y_k^H$$

where E is the cyclic downshift matrix

$$\begin{bmatrix} 0\,0\,0 \ldots 0\,1 \\ 1\,0\,0 \ldots 0\,0 \\ 0\,1\,0 \ldots 0\,0 \\ \quad\quad \ldots \\ 0\,0\,0 \ldots 1\,0 \end{bmatrix} \quad .$$

Once again, however, the computationally intensive SVD is required. In addition, α is not necessarily a small number.

4. Conclusions

For space-invariant point-spread functions and equal sampling in image and reconstruction spaces, expansion of the imaging matrix A to circulant form is straight-forward, providing immediate access to one-dimensional FFT construction techniques and greatly enhanced efficiency through reduced computational effort and storage requirements. The matrix representing the image is vectorized by a simple column-wise (or row-wise) mapping.

The penalty of a circulant expansion is loss of a smaller support in reconstruction space. However, for typical images, using, for example, a CCD array, where the point-spread function is generally of much smaller extent than the image, this is not a serious limitation.

If the point-spread function is not space-invariant, or if sampling rates differ in image and reconstruction, the SVD is the preferred means for synthesis of the deconvolution matrix. For this case, faster SVD procedures would be a valuable development. If the accuracy and reliability afforded by the SVD are not of paramount importance, fast approximate methods, based, for example, on a separable approximation to the point-spread function, may be adequate.

References

1. J. B. Abbiss, J. C. Allen, R. P. Bocker and H. J. Whitehouse, "Fast image reconstruction based on the regularized pseudoinverse of the imaging operator", in *Inverse Problems in Scattering and Imaging*, Michael A. Fiddy, Editor, Proc. SPIE 1767, pp. 93-111, 1992.

2. A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*, V. H. Winston & Sons, Washington DC, 1977.

3. A. Albert, *Regression and the Moore-Penrose Psudoinverse*, Academic Press, New York and London, 1972.

4. G. Wahba, "Practical approximate solutions to linear operator equations when the data are noisy", *SIAM J. Numer. Anal.*, Vol. 14, No. 4, pp. 651-667, 1977.

5. P. C. Hansen, "Analysis of discrete ill-posed problems by means of the L-curve", *Argonne National Laboratory Preprint* MCS-P157-0690, July 1990.

6. G. H. Golub and C. Reinsch, "Singular value decomposition and least-squares solutions", *Numer. Math.*, Vol. 14, pp. 403-420, 1970.

7. T. Kailath, S. Y. Kung and M. Morf, "Displacement ranks of matrices and linear equations", *J. Math. Anal. and Appl.*, Vol. 68, pp. 395-407, 1979.

8. J. M. Speiser, H. J. Whitehouse and J. C. Allen, "Fast matrix-vector multiplication using displacement rank approximation via an SVD", in *SVD and Signal Processing, II: Algorithms, Analysis and Applications,* Richard J. Vaccaro, Editor, Elsevier, Amsterdam and New York, 1991.

9. P. J. Davis, *Circulant Matrices*, John Wiley & Sons, New York, 1979.

10. G. Ammar and P. Gader, "New decompositions of the inverse of a Toeplitz matrix", in: A. A. Kaaschock, J. H. Van Schuppen and A. C. M. Ran (Eds.), *Signal Processing, Scattering and Operator Theory, and Numerical Methods*, Birkhauser, Boston, 1990.

APPENDIX

Related Reports Produced under Preceding Contract No. N00014-90-C-0173

1. Abbiss, J. B., Brames, B. J. and Fiddy, M. A., "Superresolution Algorithms for a Modified Hopfield Neural Network", IEEE Transactions on Signal Processing, 39, 7, 1516-1523, July 1991.

2. Abbiss, J. B. and Brames, B. J. , "Restoration of Sub-Pixel Detail Using the Regularized Pseudo-Inverse of the Imaging Operator", in *Advanced Signal Processing Algorithms, Architectures and Implementations II*, Franklin T. Luk, Editor, Proc. SPIE 1566, 365-375, 1991.

3. Abbiss, J. B., Allen, J. C., Bocker, R. P. and Whitehouse, H. J. "Fast Image Reconstruction Based on the Regularized Pseudo-Inverse of the Imaging Operator", in *Inverse Problems in Scattering and Imaging*, Michael A. Fiddy, Editor, Proc. SPIE 1767, 93-111, 1992.

Superresolution Algorithms for a Modified Hopfield Neural Network

John B. Abbiss, Bryan J. Brames, and M. A. Fiddy, *Member, IEEE*

Abstract—**The purpose of this paper is to describe the implementation of a superresolution (or spectral extrapolation) procedure on a neural network, based on the Hopfield model. This was first proposed by Abbiss *et al.* [1]. We show the computational advantages and disadvantages of such an approach for different coding schemes and for networks consisting of very simple two state elements as well as those made up of more complex nodes capable of representing a continuum. With the appropriate hardware, we show that there is a computational advantage in using the Hopfield architecture over some alternative methods for computing the same solution. We also discuss the relationship between a particular mode of operation of the neural network and the regularized Gerchberg–Papoulis algorithm.**

I. Introduction

THERE are several models of neural networks, each of which has a structure based loosely on our view of biological nervous system components [2]. A neural network architecture is one consisting of a very large number of simple processing elements densely interconnected by a set of weighted links. Each processing element updates its state by comparing the sum of its inputs with a prescribed threshold. The study of the properties of neural networks is a subject still somewhat in its infancy, and current hardware limitations reduce their practical impact. Indeed, it has been suggested by Anderson and Rosenfeld [3] that they may not become useful until inexpensive special purpose parallel hardware is available. Should that hardware be available, the question remains as to how one would make best use of a neural computer; i.e., how one should program or "train" it to perform the tasks required. The hope is that some problems for which it is difficult to find satisfactory algorithmic solutions might be amenable to this kind of computing architecture, which can organize itself and learn what it is expected to accomplish.

One anticipated use of neural networks is in autoassociative memory and in image (or signal) classification, recognition or understanding applications; these are applications that we believe the human brain is particularly good at while current algorithms implemented largely on serial machines still leave much to be desired. For most neural networks, their learning or restoration capabilities can be expressed in terms of the minimization of some appropriate energy or cost function. One of the objectives of this paper is to take an established algorithm in image reconstruction and identify those aspects of it that can be related to the programming requirements that would be necessary for implementation on a Hopfield neural computer. From the analysis of such a network applied to solve a problem for which the cost function is well defined, one might be able to assess their use for the solution of a wider class of optimization problems.

The Hopfield network is a fully connected network in the sense that any one of the processing elements is connected to every other one. This contrasts with layered networks, such as a multilayered perceptron (MLP), in which processing elements are arranged with connections only between neighboring layers. This difference in topology is accompanied by differences in the thresholding functions and in the procedures to find the connection strengths. The Hopfield network operates iteratively; the connection strengths are assigned and specify a cost function which the iterative procedure minimizes. The MLP is a one-pass network once the connection strengths have been "learned" by the minimization of an error function which quantifies the difference between the current and desired output states.

It was pointed out by Jau *et al.* [4] that some iterating image restoration processes are mathematically very similar to autoassociative memory; indeed if the input information is incomplete, it can be considered as a key pattern to an associative memory. Since the approach to image restoration presented here was first proposed [1], there have been other related studies which we mention here. Zhou *et al.* [5] considered an energy function identical to (6) in order to specify network interconnection strengths. Their application was the restoration of grey level images degraded by a shift invariant FIR blur function and additive noise. Grey level information was coded by a simple sum of binary elements and the network was serially updated with a stochastic thresholding rule to avoid getting trapped in local minima of the energy function. Jang *et al.* [6] utilized the optimization properties of the Hopfield network in order to estimate a matrix inverse. This is clearly important in image restoration problems, as indicated by (8). Full grey level representation is assumed and

Manuscript received April 25, 1989; revised August 23, 1990. This work was supported in part by the SDIO/IST and managed by ONR.

J. B. Abbiss and B. J. Brames are with Spectron Development Laboratories, Inc., Costa Mesa, CA 92626.

M. A. Fiddy is with the Department of Electrical Engineering, University of Lowell, Lowell, MA 01854.

IEEE Log Number 9144725.

1053-587X/91/0700-1516$01.00 © 1991 IEEE

a differential mode of implementation applied; they point out that this method is similar to a steepest descent method but with a nonlinear thresholding step at each iteration. Bai and Farhat [7] also considered a cost function similar to that in (6) to recover images from limited Fourier data. In their approach there was an additional constraint on the norm of the derivative of the estimate as well as on the norm of the estimate itself. The increment added to the current estimate was weighted by a "gain" factor prior to thresholding, which was chosen to ensure that the network energy function decreased at each step. Their reconstructions showed advantages for low-signal-to-noise ratio situations and are being implemented on optoelectronic hardware. Winters [8] considers a norm minimization as expressed by (6) but without any explicit regularization term included. The minimization of his cost function is achieved by the penalty method which requires that a large positive value is added to the cost function, wherever a nonlinear inequality constraint is not satisfied. An adaptive penalty function allows one to avoid local minima and this complete procedure can be mapped onto the Hopfield energy function. Results showed that the reconstructions were robust against noise and could be implemented in microseconds on an analog electronic network, as compared to several hours on a minicomputer.

Using a Hopfield network, our interest is in an application for which a solution state evolves through the minimization of some specific cost or energy function. Once the energy function is defined, one can determine the appropriate connection strengths in order that the energy function associated with the network is the same as that of the problem under consideration. A key feature of a net of this type is its construction from a set of simple processors, each of whose states is determined by a thresholding operation applied to a sum of weighted inputs from other processors or nodes. The properties of the network as a whole are determined by the thresholding function used, and by the pattern and strengths of the connections between the processing elements.

II. The Hopfield Network with Two State Elements: Theoretical Framework

The Hopfield neural model [9], [10] allows one to specify a set of desired memories as minima of a configurational energy of the network. We assume that the network consists of N processing elements each of which has two states and each of which has a thresholding operator that determines the states of the element from the total input to that element.

Given an initial starting configuration or state of the network, each processor or "neuron" updates its state according to a threshold rule of the form

$$\text{if } \sum_{j=1}^{N} T_{ij} v_j > 0 \text{ then } v_i = 1; \qquad \text{otherwise } v_i = -1 \tag{1}$$

where the elements of the connection matrix T are formed according to the rule

$$T_{ij} = \sum_{\mu=1}^{M} v_i^{\mu} v_j^{\mu} \qquad (i, j = 1, 2 \cdots N) \tag{2}$$

and the v_i^{μ} are the elements of the M memory vectors, v^{μ}, to be stored. The v_i^{μ} can take the values ± 1, and it is assumed that the diagonal term, T_{ii}, is zero in the Hopfield model.

This thresholding rule can be applied in series (asynchronously) or in parallel (synchronously). In the serial mode, the rule is applied sequentially to the nodes of the network, the state of the network being updated after each operation. In the parallel mode, the current network state remains unchanged until the thresholding operation has been applied to every node. The configurational energy function for the network has the form [9], [10]

$$E = -\tfrac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} T_{ij} v_i v_j = -\tfrac{1}{2}\, v^T T v \tag{3}$$

where superscript T denotes transpose. Serial thresholding will always minimize this energy function, provided that the T_{ii} are nonnegative. If M is sufficiently small, this state will correspond to the memory closest in Hamming distance to the state in which the network was started. Parallel thresholding results in either convergence to a stable state or oscillation between two states [11].

This iterative scheme can be expressed more concisely and modified to permit biasing of the neuron inputs in the following manner. We let the state of the network after the nth iteration be described by the N element vector

$$v^{(n+1)} = U(Tv^{(n)} + b)$$

where U is the threshold operation, T denotes the matrix with elements T_{ij}, a superscript denotes iteration number, and b is the bias vector. The bias vector incorporates boundary conditions such as image data; it effectively shifts the decision threshold for each element. In this case the energy function minimized by the network is of the form [10]

$$E = -1/2\; v^T T v - b^T v. \tag{4}$$

The change in energy for a change in the state of one neural element from v_k to $v_k + \Delta v_k$ is

$$\Delta E_k = -\Delta v_k [(Tv + b)_k + \tfrac{1}{2}\, T_{kk} \Delta v_k].$$

Taking T_{kk} to be zero ensures that the change in energy is always negative, since the term in the brackets above then has the same sign as Δv_k. If T_{kk} is nonzero, the term in the bracket will have the same sign as Δ_k provided T_{kk} is positive, and then E is guaranteed to reduce; the consequences of varying T_{kk} were explored in an earlier publication and verify the expected behavior [12]. If two (or more) neurons change state simultaneously, the change in E contains terms involving products of the form $-1/2$ $T_{kl} \Delta v_k \Delta v_l$ (or these plus higher order terms if more neurons change), the sign of which can vary.

The two state representation is too limited for a one-to-one mapping between elements and signal or image samples, in most cases. However, these simple elements can be taken in groups to represent grey levels through a variety of coding schemes. Alternatively, either analog or more complex digital processing elements could be used to directly represent a grey level.

III. The Superresolution Problem

There are many applications that require the restoration of a signal or image from a limited discrete data set; for example, samples of the spectrum of a function, or of its low or band-pass filtered image. An important *a priori* assumption for work in super resolution is the fact that most objects to be imaged are of compact support. This leads to the well-known result that their spectra are band-limited functions. In principle, therefore, one might hope to extend limited spectral data by means of analytic continuation. This procedure is notoriously unstable in the presence of noise and does not provide a practical solution to the problem. One has infinite freedom in interpolating and extrapolating limited sampled data; hence, one is forced to approach super resolution from an optimization point of view [13]. The best that one can hope to achieve is the specification of a cost or energy function which possesses a unique minimum and is designed to incorporate whatever constraints and *a priori* knowledge might be available to help limit the set of possible solutions to the problem, while retaining desirable and necessary solution characteristics. Examples of constraints include data consistency, support consistency and, perhaps, positivity. The objective of the superresolution process is to obtain a final image that has a higher spectral or spatial frequency content than the original data set, as a direct consequence of incorporating the prior knowledge available into the cost function. It is a matter of taste, to a large extent, how one designs a cost function in order to obtain a desirable solution to the problem; i.e., a superresolved signal or image with acceptable properties. The superresolution problem is thus transformed into one of determining the (global) extremum of a cost function on the assumption that this solution is optimum.

One of the early successes of a neural network was to find a good approximation to the traveling salesman optimization problem [14]. The superresolution optimization problem can be mapped onto a neural network in two distinct ways. One is to train network using a data base of superresolved images [15]–[17], the other is to relate the cost function associated with a given network to the chosen superresolution cost function. It is the latter approach that we adopt here.

IV. Superresolution and Spectral Estimation

Most signal or image recovery problems can be described by linear equations of the form

$$g(x) = \int A(x, y) f(y)\, dy$$

where A is the system spread function or the Fourier transform kernel, for example. The interpretation of the data $g(x)$ to obtain information about the object $f(y)$ requires the solution of a linear inverse problem. This is equivalent to finding the solution of a Fredholm integral equation of the first kind. It is well known that small fluctuations in the data $g(x)$ can lead to very large fluctuations in the estimate of the unknown function $f(y)$. This is a manifestation of the ill-posed nature of the problem (the inverse of the operator A is not generally continuous) and some degree of regularization is required in order to determine stable and meaningful solutions. One usually proceeds by assuming that the desired solution belongs to the space F of (possibly weighted) L^2 functions, the regularization restricting that solution to conform to any *a priori* knowledge available about the object whose enhanced image is sought.

In practice, the solution is determined from a finite set of samples of $g(x)$, and the data vector g is expressed by

$$g = Af + n$$

where A is the imaging operator and n represents an additive noise component; A explicitly contains the support constraint of f, which is assumed to be known or estimated *a priori*. These limited data can be regarded as a noisy finite set of bounded linear functionals of f.

A data-consistent solution exists, however, which is a solution of minimum norm. This solution is the data-consistent $\hat{f}$ which minimizes $\|\hat{f}\|^2$, where $\|\ \|$ denotes norm.

The solution to this minimization problem can be written

$$\hat{f} = \sum_{i=1}^{N} (g, v_i) u_i / \alpha_i \tag{5}$$

where N is the number of image data points and the α_i, u_i and v_i are the singular values, singular functions, and singular vectors, respectively, pertaining to the operator A: $Au_i = \alpha_i v_i$; $A^* v_i = \alpha_i u_i$. The singular values tend to zero as i increases, leading to an instability in the estimator. If the first $N_i \leq N$ singular values are dominant, then the remainder may be neglected, but only at the expense of loss of resolution in $\hat{f}$.

Thus, this solution is ill-conditioned but stability can be restored by relaxing data consistency; hence, we minimize the cost function

$$E = \|A\hat{f}' - g\|^2 + \beta \|\hat{f}'\|^2. \tag{6}$$

The estimate is

$$\hat{f} = \sum_{i=1}^{N} (g, v_i) u_i \alpha_i / (\alpha_i^2 + \beta) \tag{7}$$

where the regularization parameter β is chosen to achieve a compromise between resolution and stability, and usually requires some adjustment in order to establish its optimal value. As β tends to zero, the solution becomes more

data consistent. The minimizer of this cost function can be computed directly in matrix form, namely,

$$\hat{f} = [A^*A + \beta I]^{-1}A^*g \tag{8}$$

where, for a real-valued matrix, A^* becomes A^T, and I represents the identity matrix.

An alternative approach to estimating the object is to consider the minimization of the cost function $\|f - \hat{f}'\|^2$ using a trigonometric polynomial of the form [13]

$$\hat{f} = \sum_{k=1}^{N} d_k \phi_k$$

where the ϕ_k form a basis in the data space F and the optimal d_k satisfy

$$\sum_{m=1}^{N} [(\phi_m, \phi_n)_F + \beta\delta_{nm}]d_m = G_n \tag{9}$$

and the G_n are Fourier data corresponding to the low-pass filtered image g.

It is worth pointing out that in the space F that incorporates the known support constraint for the function to be restored, the three solutions given by (7)–(9) are equivalent; expression (9) can be obtained from expression (8) [18]. Each method for solution is more or less computationally the same in that each requires $\sim O(N^3)$ multiplications; this was pointed out earlier in [1].

We note that expression (7) requires on the order of CN^3 multiplications, where the overhead C is large by comparison with the other methods. However, a primary concern is the ease with which the regularization parameter β can be varied; this can be done at the cost of $O(N^2)$ multiplications for (7).

V. Implementation on a Neural Network

We will now show how a superresolution algorithm equivalent to the previously described approaches can be defined on a Hopfield neural network. Several issues must be addressed. First, it is necessary to define the connection matrix from the cost function. For some problems this cannot be accomplished without performing more calculations than are required for a more conventional solution to the problem. The latter consideration was noted by Takeda and Goodman [19]. Thus, the computational load or complexity must be considered in deciding the merits of a neural network solution to this problem.

Other issues which must be addressed center on modifications of Hopfield's formulation to satisfy our requirements. Hopfield ensures that the network will converge by arbitrarily setting the diagonal of the connection matrix to zero. This is unacceptable, because it shifts the absolute energy minimum of the network from the minimum of (6). Moreover, while a suitable network can be constructed from two-state elements, one often requires that the reconstruction be represented over at least a set of grey levels. Thus, it may be necessary to combine a number of neural elements to represent a reconstruction pixel. The method of coding the grey levels depends on the complexity of the elements, i.e., whether they can only take on two states, a bounded continuum of values such as $0 \leq v_k \leq 1$ ("graded neurons"), or an (effectively) unbounded continuum. Granularity of the representation will affect the convergence properties of the network; coarsely quantized systems can converge to local energy minima, yielding less than optimal reconstructions.

We shall first demonstrate the formal mapping of a superresolution algorithm onto a Hopfield network. The technique will then be extended to fully address the second and third issues on a two-state network. We will briefly examine the advantages of more finely quantized systems, and finally discuss the relation between parallel thresholding and a regularized form of the well-known Gerchberg–Papoulis spectral extrapolation algorithm [20]–[22].

Let us represent the current estimate by the state of the network v. We can rewrite (6) as

$$E = v^TA^TAv - 2v^TA^Tg + g^Tg + \beta v^Tv.$$

Comparing this expression for E with that of the Hopfield network, (4), gives

$$\left.\begin{aligned} T &= -2(A^TA + \beta I) \\ b &= 2A^Tg \end{aligned}\right\} \tag{10}$$

where I is the identity matrix, and the g^Tg term can be ignored, since it represents a total offset for E.

Thus, superresolution can be mapped simply and directly onto a Hopfield network. The connection matrix is formed from the imaging operator matrix, which contributes information about the imaging system, and the regularization parameter β, which sets a bound on the norm of the final estimate. The available data g contribute only to the bias vector b.

For serial operation, the change in energy due to a change Δv_k is

$$\Delta E_k = -\Delta v_k[(Tv + b)_k + \tfrac{1}{2} T_{kk}\Delta v_k]. \tag{11}$$

Convergence to a minimum is guaranteed if the expression in braces always has the same sign as Δv_k. This can be ensured by altering the diagonal of T to zero; however, such a change is equivalent to choosing an arbitrary value for β. This is not acceptable, since the regularization parameter should be chosen to reflect the noise in the data and to obtain an optimum reconstruction, not to ensure convergence of the algorithm.

VI. Superresolution on a Two-State Network

In this section we will modify the Hopfield formulation so that the energy minimum of the network will coincide with that of (6), while still decreasing the energy with each change in the state vector. We shall find that this is possible by introducing a two-level threshold in place of the usual single-level one. In addition, we will incorporate a generalized grey scale mapping which describes linear transformations of a state vector v into a vector w having grey levels. We write this as $w = Sv$, where S could

be a mapping from an N element vector, each of whose elements can take the values 0 or 1, to an L-element vector ($L \leq N$) whose elements can take a wider range of values. For example, if S represents a base-2 mapping, each element of v can represent a power of 2, giving a range of $2^{N/L}$ values for each element of w. The range of v need not be limited to $\{0, 1\}$: we use this for the purpose of illustration. Other coding schemes are possible, such as clustering or bit-density coding.

The expression for the energy is now

$$E = \|Aw - g\|^2 + \beta \|w\|^2.$$

Suppose the grey-level vector w is perturbed by some amount Δw. The difference in energy between states is

$$\Delta E = 2\Delta w^T\{(A^TA + \beta I)w - A^Tg + \tfrac{1}{2}(A^TA + \beta I)\Delta w\} \tag{12}$$

or, in terms of the neural state vector v,

$$\Delta E = 2\Delta v^T[S^T(A^TA + \beta I)Sv - S^TA^Tg + \tfrac{1}{2} S^T(A^TA + \beta I)S\Delta v]. \tag{13}$$

It should be emphasized that (13) contains no assumption about the range of values of v or of the updating mode of the network. A restriction to two states reflects the desire to use a large number of simple binary processing elements in neural architectures.

We now present a procedure which ensures that the change in energy expressed by (12) and (13) always decreases, provided serial thresholding is adopted. For a change Δv_k, the change in energy ΔE_k of the network is given by (11), with the following definitions for T and b:

$$T = -2S^T(A^TA + \beta I)S$$

$$b = 2S^TA^Tg.$$

The grey-scale mappings we are considering associate a specific neuron with one and only one image pixel. Hence, the columns of S each contain only one element, and it is not difficult to show that the diagonal elements of T take the form

$$T_{kk} = -2S_{jk}^2(A^TA + \beta I)_{jj} \tag{14}$$

where S_{jk} is the nonzero element of the kth column of S. Since the diagonal elements of A^TA are positive, and β is some positive quantity, T_{kk} is always negative. Hence, we can rewrite (13) in the form

$$\Delta E_k = -\Delta v_k\{(Tv + b)_k - \tfrac{1}{2}|T_{kk}|\Delta v_k\}. \tag{15}$$

Thus, ΔE_k will be negative provided $(\Delta v_k)^2 < \Lambda_k \Delta v_k$, where

$$\Lambda_k = \frac{2}{|T_{kk}|}(Tv + b)_k$$

the maximum decrease in energy occurring when

$$\Delta v_k = \tfrac{1}{2}\Lambda_k. \tag{16}$$

Thus, we require that $|\Delta v_k| < |\Lambda_k|$ and sgn (Δv_k) = sgn (Λ_k). For a binary network, where $v_k \in \{0, 1\}$, we then obtain the following rule to ensure that the network energy does not increase:

$$v_k^{(n+1)} = \begin{cases} 1 & \text{for } \Lambda_k > 1 \\ v_k^{(n)} & \text{for } |\Lambda_k| \leq 1 \\ 0 & \text{for } \Lambda_k < -1. \end{cases}$$

We consider next the operation of a nonbinary network.

VII. Superresolution on a Nonbinary Network

The restriction of the state vector v to binary values permitted the simplest possible processing elements to be used in a neural architecture. With more complex processors this simple representation is unnecessary and inefficient: the optimal coding scheme is intimately related to the nature of the available hardware.

A disadvantage of simple two-level elements is that they can give a coarsely quantized representation in reconstruction space which leads to the creation of local energy minima. There is still only one absolute energy minimum, but the network may converge to a local minimum of higher energy. It should be recognized that this behavior also occurs with a single level threshold and a zero-diagonal T matrix, and the resulting reconstructions are sometimes called "spurious stored states." A typical solution of this type is shown in Fig. 1(c) for a network of 90 two-level elements; a 6-b coding scheme yields 15 points in the reconstruction. Whether such a reconstruction is of acceptable quality is difficult to predict, and a function of the needs of the user. This difficulty can be overcome by using elements which can take on values over a continuum, such as $0 \leq v_k \leq 1$ (Fig. 1(e), (f)). These elements are similar to the graded neurons employed by Hopfield in a differential network [10].

We shall now examine the behavior of an asynchronous network composed of elements which can take on a continuum of values. Because $|\Delta v_k|$ is no longer fixed there is no need for a threshold/decision operator; we will simply use the value of Δv_k which yields the greatest decrease in energy.

It was noted above that the maximum decrease in energy occurs when

$$\Delta v_k = \Lambda_k/2.$$

In addition, $|\Lambda_k|$ represents an upper bound on $|\Delta v_k|$. As one approaches the solution, $(Tv + b)_k$ approaches zero, so this upper bound decreases. For networks with a fixed $|\Delta v_k|$ (e.g., ± 1 for all k), one would expect $|\Lambda_k|$ eventually to be smaller than $|\Delta v_k|$, so no changes can be made to reduce the energy, even though the network is not yet at the global minimum. Thus, the fixed step methods will generally be limited to some outer neighborhood of the global energy minimum (although one might arrive at the minimum).

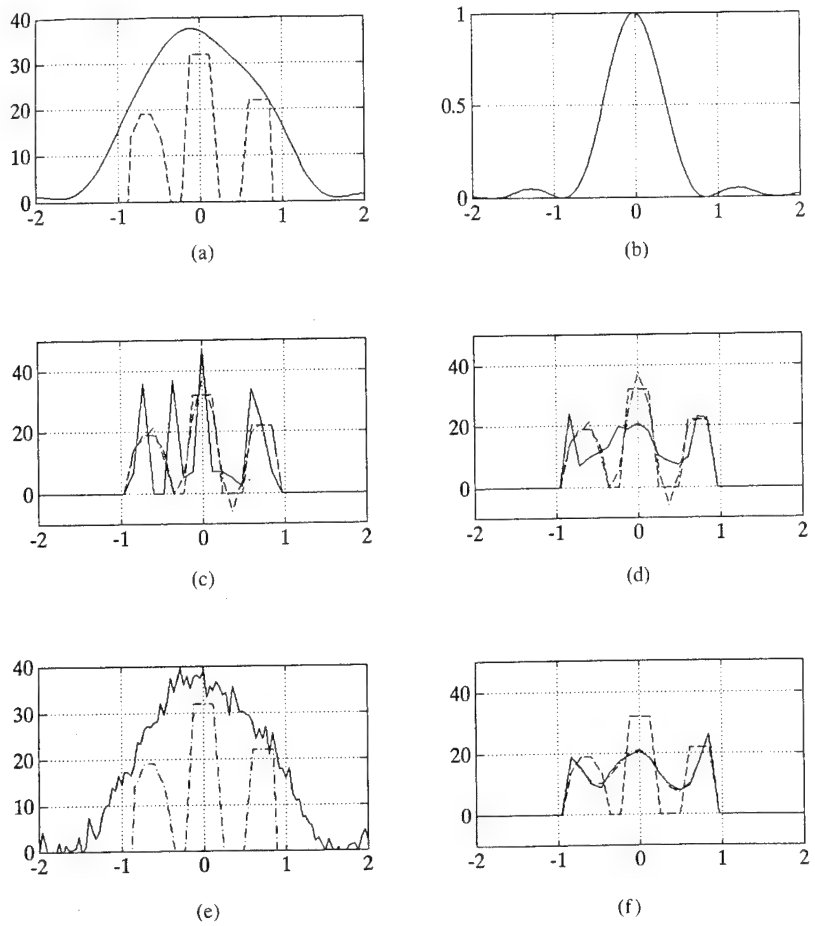

Fig. 1. The line types used are solid for an image or a neural network reconstruction; dashed for the object; and dot-dashed for the algebraic (SVD) reconstruction. (a) Object and incoherent image. Ninety two-state {0, 1} elements are mapped with a 6-b base 2 scheme to 15 pixels in the object. (b) The point-spread function of the imaging system (sinc^2). (c) Network using two-state neurons converges to a local energy minimum after 5 cycles, $\beta = 10^{-5}$. (d) After 5 cycles using 90 graded neurons the network has not converged, but a good estimate of the object has emerged, $\beta = 10^{-5}$. (e) Object and image containing 5% Gaussian additive noise. (f) Neural reconstruction from (e) after 50 cycles, $\beta = 10^{-4.5}$. Note that it is nearly indistinguishable from the SVD result.

However, if we adopt graded neurons it is still possible to use very simple elements, yet circumvent the finite step limitation. Since these elements can take on a continuum of values between two limits, the energy is forced to decrease at each step. A serially threshold network, constructed from these elements can therefore reach the global energy minimum after a sufficient number of cycles. One could also dispense with the coding scheme, and use a smaller number of more complex elements.

We would like to operate the network in the synchronous mode to make efficient use of the network's parallelism. If the kth neuron changes by Δw_k

$$w_k^{(n+1)} = w_k^{(n)} + \Delta w_k$$

where convergence is assured provided that the conditions of Section VI are met.

A regularized form of the Gerchberg–Papoulis algorithm reads [20]–[22]

$$\begin{aligned} w^{(n+1)} &= A^T g + [(1-\beta)I - A^T A]w^{(n)} \\ &= w^{(n)} + (Tw^{(n)} + b). \end{aligned}$$

Thus, if

$$\Delta w_k = (Tw^{(n)} + b)_k$$

parallel operation of the network will result in a computation which is identical to the regularized Gerchberg–Papoulis algorithm. Since the latter always converges [22], this choice for the Δw_k must always be possible. Optimal selection of the Δw_k to accelerate convergence of the network in the parallel mode is under investigation.

VIII. Computational Complexity of Neural Algorithm

The computational complexity associated with image reconstruction or superresolution using the singular value decomposition of A to solve (8), and using the neural network approach, has been examined for one-dimensional images (Table I). The computational load associated with the neural network is independent of whether the thresholding is serial or parallel, although the actual computational time is obviously less for parallel thresholding. The

TABLE I
THE NUMBER OF OPERATIONS REQUIRED FOR SUPERRESOLUTION BY SVD INVERSION AND BY A NEURAL OPTIMIZATION ARE LISTED FOR THE TOTAL CALCULATION, AND FOR UPDATING EITHER THE IMAGE OR THE REGULARIZATION PARAMETER

Requirement	Operation	SVD	Neural
New β	Mults	$N^2 + N$	$KN^2 + (K + 1)N$
	Adds	N^2	$KN^2 + (K + 1)N$
	Divs	N	N
New image	Mults	$2N^2 + N$	$(K + 1)N^2 + (K + 1)N$
	Adds	$2N^2$	$(K + 1)N^2 + KN$
	Divs	N	N
Total operations	Mults	$15N^3 + 3N^2 + 2N$	$N^3 + (K + 1)N^2 + (K + 1)N$
	Adds	N^2	$N^3 + KN^2 + KN$
	Divs	N	N

number of additions, multiplications, and divisions for each technique are listed in Table I for three situations. The first row gives the number of calculations needed for a new value of the regularization parameter β; the neural network has a disadvantage in this case because it must generally run for some K iterations. The second row gives the computational cost involved in updating the input image data vector g. The neural network once again is at somewhat of a disadvantage. However, by examining the total number of operations from the beginning, one can see that the neural approach is substantially more efficient because it calculates a matrix product once without the overhead associated with singular value decomposition; numerical experiments using a microcomputer indicate that the number of operations grows as $15N^3$ for an N point image. This is clearly increasingly significant for larger images.

IX. DISCUSSION AND CONCLUSIONS

Neural network solutions to image restoration problems are, therefore, competitive with, but not necessarily better than, more traditional methods for solving the problem (see also [22]). We have shown that both binary and nonbinary image reconstruction algorithms can be implemented on very similar neural architectures. The nonbinary case can be based upon network elements which take discrete (typically two-state) or continuous values. In the present application, the diagonal of the connection matrix is nonzero and, for the discrete element case, the crucial thresholding step was modified in order to ensure that the energy of the network decreases at each step; the thresholding step is unnecessary for the continuous case.

Thus, convergence to a minimum of the energy function is guaranteed for the both the discrete and continuous-element networks. However, this minimum is unique only in the continuous case, since discretization of the element states introduces local minima. Convergence has been demonstrated in this paper only for the case of serial thresholding; for parallel thresholding in the continuous case, the computation reduces to the regularized Gerchberg–Papoulis algorithm for a particular choice of the increments in the states of the network elements.

ACKNOWLEDGMENT

The authors acknowledge useful discussions with J. S. Bayley and C. L. Byrne.

REFERENCES

[1] J. B. Abbiss, J. S. Bayley, B. J. Brames, and M. A. Fiddy, "Superresolution and neural computing," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 880, pp. 100–106, Jan. 1988.
[2] J. A. Reggia and G. G. Sutton, "Self-processing networks and their biomedical implications," *Proc. IEEE*, vol. 76, pp. 680–692, 1988.
[3] J. A. Anderson and E. Rosenfeld, Eds., *Neurocomputing: Foundations of Research*. Cambridge, MA: M.I.T. Press, 1987.
[4] J. Y. Jau, Y. Fainman, and S. H. Lee, "Comparison of adaptive pattern recognition and image restoration with heteroassociative and autoassociative memories," in *Tech. Dig. O.S.A. Topical Meeting Opt. Comput.*, 1987, pp. 145–148.
[5] Y.-T. Zhou, R. Chellappa, A. Vaid, and B. K. Jenkins, "Image restoration using a neural network," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 1141–1151, 1988.
[6] J.-S. Jang, S.-Y. Lee, and S.-Y. Shin, "An optimization network for matrix inversion," in *Neural Information Processing Systems*, D. Z. Anderson, Ed. AIP, 1988, pp. 397–401.
[7] B. Bai and N. H. Farhat, "Radar image reconstruction based on neural net models," in *Proc. IEEE APS/URSI Meeting* (Syracuse, NY), 1988, pp. 774–777.
[8] J. H. Winters, "Superresolution for ultrasonic imaging in air using neural networks," in *Proc. IEEE Int. Conf. Neural Networks* (San Diego, CA), 1988, pp. 609–611.
[9] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Nat. Acad. Sci. U.S.*, vol. 78, pp. 2554–2558, 1982.
[10] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two state neurons," *Proc. Nat. Acad. Sci. U.S.*, vol. 81, pp. 3088–8092, 1984.
[11] J. Bruck and J. W. Goodman, "A generalized convergence theorem for neural networks and its applications in combinational optimization," in *Proc. IEEE Int. Conf. Neural Networks*, 1987, pp. 649–656.
[12] J. S. Bayley and M. A. Fiddy, "On the use of the Hopfield model for optical pattern recognition," *Opt. Commun.*, vol. 64, pp. 105–110, 1987.
[13] C. L. Byrne, R. M. Fitzgerald, M. A. Fiddy, A. M. Darling, and T. J. Hall, "Image restoration and enhancement," *J. Opt. Soc. Amer.*, vol. 73, pp. 1481–1487, 1983.
[14] J. J. Hopfield and D. W. Tank, "Neural computation of decisions in optimization problems," *Biolog. Cybern.*, vol. 52, pp. 141–152, 1985.
[15] N. H. Farhat and S. Miyahara, "Superresolution and signal recovery using models of neural networks," *OSA Top. Meeting Signal Recovery and Synthesis II*, 1986, pp. 120–123.
[16] R. Rastogi, P. K. Gupta, and R. Kumaresan, "Array signal processing with interconnected neuron-like elements," in *Proc. ICASSP*, 1987, pap. 54.8.1, pp. 2328–2331.

[17] G. Eichmann and M. Stojancic, "Superresolving signal and image restoration using a linear associative memory," *Appl. Opt.*, vol. 26, pp. 1911–1918, 1987.
[18] A. M. Darling, T. J. Hall, and M. A. Fiddy, "Stable, noniterative object reconstruction from incomplete data using *a priori* data," *J. Opt. Soc. Amer.*, vol. 73, pp. 1466–1469, 1983.
[19] M. Takeda and J. W. Goodman, "Neural networks for computation: Number representations and programming complexity," *Appl. Opt.*, vol. 25, pp. 3033–3046, 1986.
[20] A. Papoulis, "A new algorithm in spectral analysis and band-limited extrapolation," *IEEE Trans. Circuits Syst.*, vol. CAS-22, pp. 735–742, 1975.
[21] R. W. Gerchberg, "Superresolution through error reduction," *Opt. Acta*, vol. 21, pp. 709–720, 1974.
[22] J. B. Abbiss, C. DeMol, and H. S. Dhadwal, "Regularized iterative and noniterative procedures for object restoration from experimental data," *Opt. Acta*, vol. 30, pp. 107–124, 1983.

*

John B. Abbiss received the B.Sc. degrees in mathematics and physics and the M.Sc. degree in solid state physics from the University of Wales, U.K., and later the Ph.D. degree in photon correlation techniques from the University of Surrey, U.K.

He spent the early part of his career working on guided missile development, after which he joined the Royal Aircraft Establishment at Farnborough. His research activities there involved probabilistic modeling for missile trials, determination of particle-size distributions from scattered-light measurements, and the development of photon correlation laser anemometry techniques for wind-tunnel applications. Other research interests included the development of improved methods for spectral estimation and band-limited signal extrapolation, and for image restoration from blurred and noisy data. Since joining Spectron (a division of the Titan Corporation) in California in 1985, he has been engaged in a wide range of projects, including the design and development of signal processing methods, software and instrumentation for laser velocimetry and optical spectroscopy in ground-based and airborne sensor systems, the further development of procedures for signal and image recovery, and investigations of applications of these techniques to satellite surveillance, together with their implementation in neural network form.

Bryan J. Brames was born in Ft. Wayne, IN. He received the bachelor's degree in physics from Purdue University, and the Ph.D. degree in optics from the University of Rochester in 1985.

He conducted postdoctoral work at the Blackett Laboratory, Imperial College. He then joined Spectron Development Laboratories in 1987, where he has participated in the development of optical air-data sensors, solid-state lasers, and signal processing, including phase retrieval and superresolution.

*

M. A. Fiddy (M'88) was born in Portsmouth, England, and attended the University of London. He received the first class B.Sc. degree in physics in 1973, and the Ph.D. degree in 1977.

He worked as a postdoctoral Research Assistant in the Physics Department at Queen Elizabeth College, University of London in 1977, and in the Department of Electronic and Electrical Engineering at the University College London in 1978. In 1979 he was appointed a Lecturer in Physics at Queen Elizabeth College and then moved to Kings College London in 1983. During part of 1982, he was a Visiting Associate Professor at the Institute of Optics. University of Rochester and in 1985–1986 he held a similar position in the Mathematics Department at the Catholic University of America, Washington, DC. In September 1987 he joined the faculty of the University of Lowell. He has worked for many years in the field of inverse problems and optical signal processing; he has written over 60 refereed publications and over 90 conference publications.

Dr. Fiddy is on the Editorial Board of the journals *Inverse Problems* (Institute of Physics), *Waves in Random Media* (Institute of Physics), *Multidimensional Systems and Signal Processing* (Kluwer), and *Optical Computing and Processing* (Taylor and Francis). He is also a member of OSA, IOP (U.K.), and SPIE.

PROCEEDINGS REPRINT

SPIE–The International Society for Optical Engineering

Reprinted from

Advanced Signal Processing Algorithms, Architectures, and Implementations II

24–26 July 1991
San Diego, California

Volume 1566

Restoration of sub-pixel detail using the regularized pseudo-inverse of the imaging operator

J. B. Abbiss and B. J. Brames

TITAN SPECTRON Division
3535 Hyland Avenue
Costa Mesa, California 92626

ABSTRACT

We present an analysis and computational results relating to the regularized restoration of subpixel information from undersampled data. The method makes use of a small set of images in various stages of defocus. An iterative implementation permits the incorporation of a non-negativity constraint. The problem we consider is fundamentally under-determined, but useful results can be obtained in reasonably low noise conditions.

1. INTRODUCTION

The investigations discussed here form part of a program whose subject is the enhancement of images obtained from space-based remote sensors. For the present purpose, these images are assumed to consist of quantized data from a fixed two-dimensional set of sensors, such as a CCD array. Typically for these arrays, the Airy disc is smaller than one pixel. Thus, for reasonably fast and well-corrected optics, the conventional limit on system resolution is likely to be the result of the spatial sampling associated with the pixel size; i.e., the integration of the light energy in the image over the area represented by each pixel.

For clarity, we consider the case of one-dimensional imaging with an incoherent source. Let the system point spread function (psf) be represented by the continuous imaging operator L . Then, for an isoplanatic system in the absence of noise, the image g of an object f is given by the convolution

$$g(y) = \int_S L(x - y)\, f(x)\, dx \tag{1}$$

where S is the support of f . The output of the k^{th} detector (pixel), extending from Y_k to Y_{k+1} , is

$$g_k = \int_{Y_k}^{Y_{k+1}} g(y)\, dy$$

$$= \int_{Y_k}^{Y_{k+1}} dy \int_S L(x - y)\, f(x)\, dx$$

Interchanging the order of integration (which is certainly permissible with the physical functions considered here) and making the definition

$$a_k(x) = \int_{Y_k}^{Y_{k+1}} L(x - y)\, dy \ ,$$

we obtain for the pixel-integrated image

$$g_k = \int_S a_k(x)\ f(x)\ dx,\ k = 1, 2, \cdots, K \tag{2}$$

We now discretize the object into a set $\{f_j, j = 1, 2 \cdots, J\}$ over equal intervals and write, as an approximation,

$$g_k = \sum_{j=1}^{J} a_{kj} f_j, \quad k = 1, 2, \cdots, K \tag{3}$$

where a_{kj} is the integral of $a_k(x)$ over the j^{th} interval. In matrix form, equation (3) can be written

$$g = Af$$

where A is a K x J matrix.

Since the a_{kj} can be determined from the system psf and the pixel array structure, the deconvolution problem represented by equation (3) can, in principle, be solved and the set $\{f_j\}$ reconstructed approximately from an equal set of measurements of the g_k.

In practice, the restoration problem is made much more difficult because of background and intrinsic noise, array imperfections and errors or uncertainties in the knowledge of the optical properties of the system. When pixel integration is over a significant part of the system psf, achieving even a modest degree of enhancement, using a single image, poses intractable difficulties in the presence of these perturbations. However, the necessary additional information can be derived from multiple differing images of the given object. After a brief discussion of the characteristics of the inverse problem represented by equation (3), we shall give a detailed description of a specific method for acquiring this information.

2. REGULARIZED IMAGE RESTORATION

There are many applications that require the restoration of a signal or image from a limited discrete data set; for example, samples of the spatial or temporal spectrum, or of the object's low or band-pass filtered image. An important *a priori* assumption in image restoration is that the object or objects are of compact support. This leads to the well-known result that their spectra are bandlimited functions. In principle, therefore, one might hope to extend limited spectral data by means of analytic continuation, and then, by Fourier transformation, obtain an enhanced image. This procedure is notoriously unstable in the presence of noise and does not provide a practical solution to the problem.

This central difficulty can be expressed in another way. Most signal or image recovery problems can be described by linear equations of the form

$$g(x) = \int A(x,y)\, f(y)\, dy$$

where A is the system point spread function or the Fourier transform kernel, for example. The interpretation of the data g(x) to obtain information about the object f(y) requires the solution of a linear inverse problem. This is equivalent to finding the solution of a Fredholm integral equation of the first kind. It is well-known in such a case that small fluctuations in the data g(x) can lead to very large fluctuations in the estimate of the unknown function f(y). This is a manifestation of the ill-posed nature of the problem (the inverse operator is unbounded) and some method of stabilization is needed to determine useful solutions.

The problem of a lack of continuous dependence on the data can be overcome by one of the various techniques of regularization [1, 2]. Essentially, the ill-posed problem is replaced by a related well-posed one, chosen to be physically meaningful and to possess the necessary properties of convergence and stability. Thus we may change the concept of a solution, or the Hilbert spaces of which the object and image are elements, or their topologies, or the operator itself. The technique we shall use belongs to the last category.

We impose physically reasonable constraints on the permitted solutions. If ε is a measure in the norm sense of the noise in the image (norm in the appropriate Hilbert space is denoted by $\| \cdot \|_H$) and if C is some constraint operator with E a known bound, we shall require that all possible reconstructions f' satisfy

$$||g - Af||_G \leq \varepsilon \quad \text{and} \quad ||Cf||_F \leq E$$

C may be used, for example, to impose smoothness on the reconstruction or to weight the reconstruction support. If C is the identity operator, E is a bound on the norm of the reconstruction. We combine the constraints quadratically and minimize the functional

$$||g - Af||_G^2 + \beta ||Cf||_F^2$$

where $\beta = \varepsilon^2/E^2$. Note that smaller values of β are equivalent to demanding greater fidelity between the reconstruction and the data; greater values place more emphasis on the property of the reconstruction associated with C. In the present discussion we shall take C = I. The minimizer f_β can then be expressed in either of the forms

$$f_\beta = (A^*A + \beta I)^{-1} A^* g$$

or (4)

$$f_\beta = A^* (AA^* + \beta I)^{-1} g$$

where A* is the operator adjoint to A. The inverses of the bracketed operators will always exist, since the eigenvalues of the symmetric operators A*A and AA* are non-negative. The operator, the image and the reconstruction will consist in practice of finite arrays and equation (4) becomes a matrix equation.

3. SUB-PIXEL RESOLUTION FROM MULTIPLE DEFOCUSSED IMAGES

We can state the deconvolution problem in the more general case where we are given a set of R differing noisy images of the same object over the same pixel array.

Then the r^{th} image consists of the vector of pixel outputs given by the equation

$$g^{(r)} = A_r f + n^{(r)} \tag{5}$$

where $n^{(r)}$ represents some additive noise vector. (Other forms of noise can be accommodated by appropriate modifications of equations (4) and (5), but we shall not address this question here.) We are required to estimate f from the set $\{g^{(r)},\ r = 1, 2, \cdots, R\}$.

We obtain the solution by assembling the image set into one composite image, and the corresponding matrices of integrated psf samples $a_{kj}^{(r)}$ into a single imaging matrix. The regularized solution is then again given by equation (4). The appropriate value for the regularization parameter β, which is closely related to the signal-to-noise ratio in the data, can be estimated in several ways; for example, by the method of weighted cross-validation [3].

The image set required to achieve sub-pixel resolution can be derived by various means. Stark and Oskoui [4] discuss an object reconstruction technique which uses a set of images differing from one another by rotation or lateral translation of the pixel array. It is evident that acquiring a sequence of data sets by lateral displacement of the detector (or image) through some fraction of a pixel at each step allows one to sample the image as finely as is desired. (For the two-dimensional image, the displacements can be in any direction.) Rotational displacement similarly permits arbitrarily fine sampling. They consider the specific case in which the system psf is much smaller than a pixel, so that resolution in the data is governed almost completely by pixel integration rather than the optical properties of the system.

An alternative method, which we shall adopt, is to use an image set obtained with differing point-spread functions. These could be generated by separate optical systems, or conceivably at different wavelengths. The system psf can also be conveniently altered by varying the degree of defocus; implementations of this method could depend on a single detector translated into the chosen planes of defocus, or a system of beamsplitters and detector arrays in appropriate locations.

If the reconstruction procedure is to be effective, the various images must contain significantly different information, which implies that the point spread functions must differ appreciably over the scale of a pixel. We shall assume that the images are formed on the same array, or arrays with identical characteristics, and that the object field is spatially and, if necessary,

temporally invariant from one image to another. We shall also assume that the point spread functions are accurately known and, for the purposes of the algorithm used later, that the images are formed from incoherent radiation. We do not require that any of the point spread functions are much smaller than a pixel; for the illustration presented below, the psf at focus is about half the size of a pixel. In a practical implementation, the images would not be centered on the pixel arrays in exactly the same way; i.e., there would be some lateral translation between the images. The effect of including lateral translations has not yet been investigated, but one would not expect reconstruction quality to suffer under these circumstances.

4. ILLUSTRATION

To illustrate these ideas, we include some results from a numerical simulation. The object field consists of a group of delta-function-like incoherent radiators. We shall show that from a small set of images, one at focus, the others at various stages of defocus, object locations can be well recovered, even when several of the objects' geometrical images lie within a single pixel. The algorithm is based on the regularized formula of equation (4); in addition, the calculation is iterated a small number of times to enforce non-negativity on the reconstruction.

The reconstruction is made initially into a spatial region defined by the central lobe of the focussed image, but using a finer grid. No prior assumptions are made about the locations or the number of objects within this region. For small, relatively isolated sources, some ringing will occur in the reconstruction, with associated negative pixel values. The support is progressively refined by eliminating these pixels at each iteration until an entirely positive reconstruction is obtained. The smallest object space which is consistent with the image data will yield the best reconstruction, and the problem, initially underdetermined, becomes finally an overdetermined one.

A modified form of this scheme, which would be appropriate for more extended objects, includes a weight matrix which biasses the next iteration against those pixels with negative values. The computation is significantly slower in this case, since the size of the reconstruction space remains constant. We also note the possibility of using a regularized form of the non-negative least-squares algorithm of Lawson and Hanson [5]. The relative performance of this procedure, also of course iterative, has not been fully evaluated.

In this example, the object field consists of eleven highly localized sources of equal intensity, distributed over a 3x3 block of image pixels; see Figure 1. (It should be noted that the reconstruction grid used did not coincide with the object grid; hence the objects cannot appear as single-pixel "points" in the reconstruction.) The central lobe of the system psf was about half the width of a pixel. Four independent images containing equal energy were generated over a 7x7 block of pixels, the first corresponding to a focussed system, the others at various stages of defocus. A method was devised for choosing defocus conditions with significantly different information content. Starting with the focussed image, with associated imaging matrix A_0, we wish to find a defocussed image whose imaging matrix A_1 is as independent as possible from A_0. The criterion used for this purpose was the magnitude of the condition number (the ratio of the largest to the smallest singular values) of the matrix formed from the center columns of A_0 and A_1. The range of defocus over which the search was made was from 0 to 4 waves. The combination selected was that possessing the smallest condition number. Knowing A_0 and A_1, A_2 was determined, and finally A_3. The merit functions (reciprocals of condition numbers) calculated for combinations of two, three and four defocus levels are shown in Figure 2. Note that a zero occurs whenever the variable degree of defocus coincides with that corresponding to one of the other imaging matrices; then the composite matrix is singular and its condition number bcomes unbounded. The degrees of defocus chosen for the four images in this case were of magnitude 0, 0.8, 1.6 and 2.4 waves. The corresponding images are shown in Figures 3-6. The only readily-identifiable feature in the focussed image is that the central pixel is brighter than the surrounding ones.

Reconstructions were performed over the region defined by the central block of 3x3 pixels, with sampling seven times as fine as that in the data. Thus the reconstruction space initially consisted of 441 points, while the four images provided a total of 196 data points. The reconstructions obtained when each of the images was corrupted by additive Gaussian noise with standard deviation equal to 1% of the mean pixel content is shown in Figure 7. All of the objects are located close to their true subpixel positions. Figure 8 shows the result obtained when the noise level is increased to 5%. The reconstruction is still generally accurate, although there is now some distortion in object location and one or two small artifacts have begun to appear. A signal-to-noise ratio can be defined as the ratio for each image of the sum, on a pixel-by-pixel basis, of the signal power to the sum of the noise power. At the 5% level, this quantity varied between 33 and 28 dB. This example was designed

to be reasonably challenging, and spreading the objects further apart, or reducing their number, considerably increases the algorithm's robustness against noise.

5. CONCLUSIONS

A method for recovering detail at the sub-pixel level from a small set of images in various stages of defocus has been discussed and demonstrated, using an algorithm which incorporates regularization to counter the destabilizing effects of noise. The iterative version of this algorithm, which permits the inclusion of a non-negativity constraint, is particularly effective at recovering accurate object support estimates. It is therefore an appropriate technique to use when it is known, *a priori*, that the object field consists largely of small well-separated targets. The sensitivity to noise of the method deserves more detailed investigation. A quantitative comparison of its performance with methods which make use of laterally translated or rotated image sets would also be of considerable interest. In practice there would inevitably be some lateral shift between images, even if they are formed on the same array, and a hybrid scheme might prove to be the most robust in the presence of noise. Ultimately, of course, the performance of any restoration algorithm must be limited by the information content of the image set. What should be sought, therefore, is the encoding scheme which most efficiently exploits the total information carried by the incident radiation.

6. ACKNOWLEDGEMENT

This work was supported by SDIO/IST and managed by ONR.

7. REFERENCES

1. A. N. Tikhonov, and V. Y. Arsenin, Solutions of Ill-Posed Problems, V. H. Winston and Sons, Washington, D.C. (1977).
2. K. Miller "Least squares methods for ill-posed problems with a prescribed bound," SIAM J. Math. Anal., 1, 1, 52-74 (1970).
3. G. Wahba, "Practical approximate solutions to linear operator equations when the data are noisy," SIAM J. Numer. Anal. 14, 4, 651-667 (1977).
4. H. Stark, and P. Oskoui, "High-resolution image recovery from low resolution detector arrays," Digital Image Synthesis and Inverse Optics, Arthur F. Gnitro, Paul S. Idell, Ivan J. LaHaie, Editors, Proc. SPIE 1351, 80-84 (1990).
5. C. L. Lawson and R. J. Hanson, Solving Least Squares Problems, Prentice-Hall, New Jersey (1974).

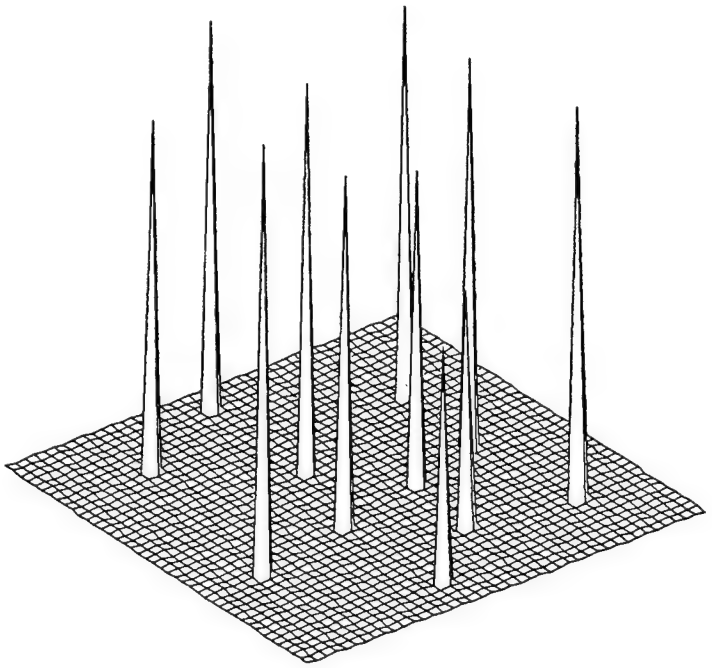
Object field

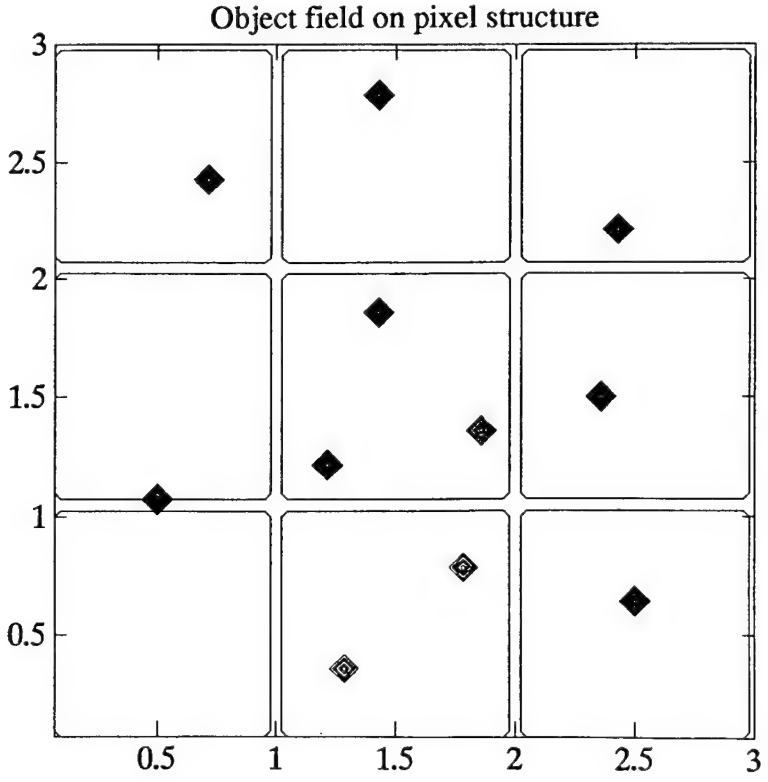
Object field on pixel structure
3
2.5
2
1.5
1
0.5
0.5
1
1.5
2
2.5
3

Fig. 1.

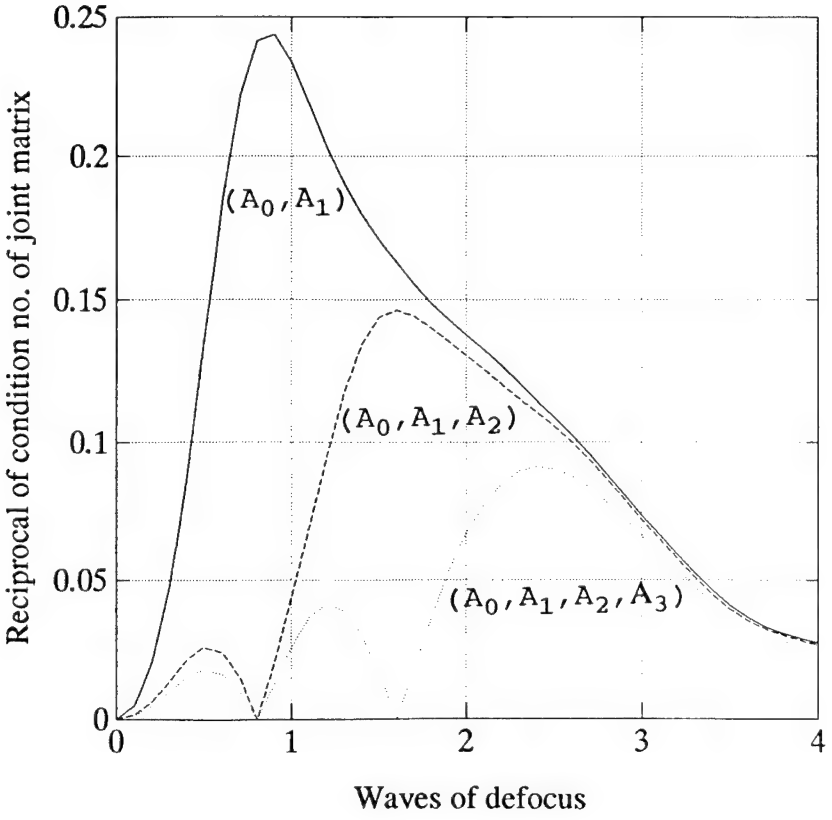

Fig. 2. Merit functions

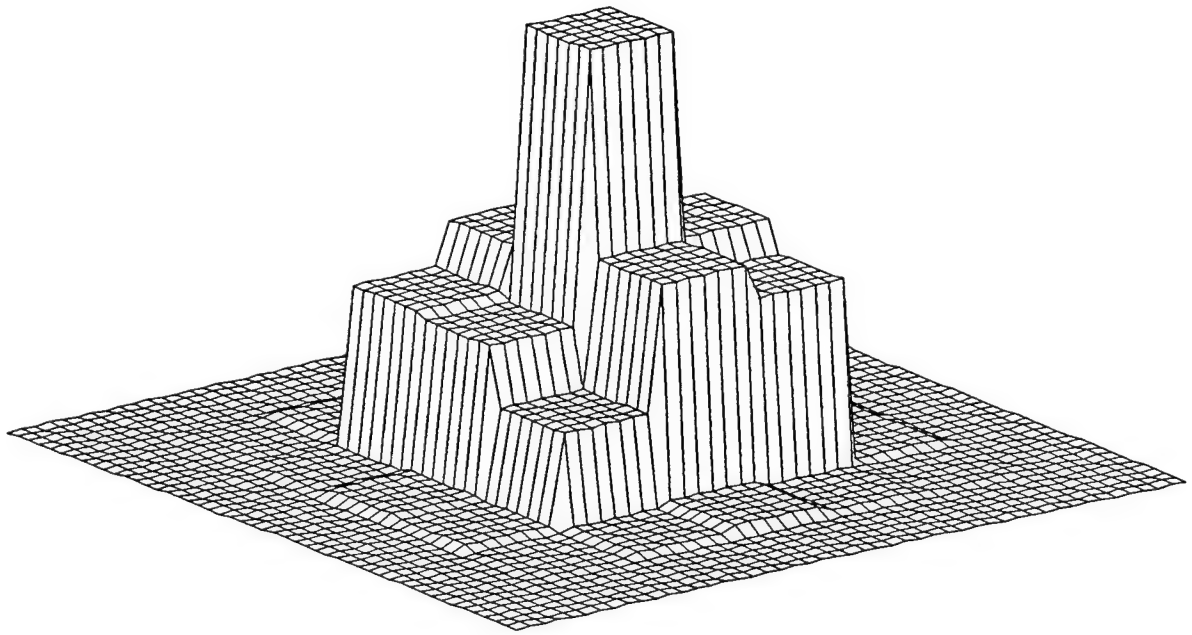

Fig. 3. Image at 0.0 waves defocus

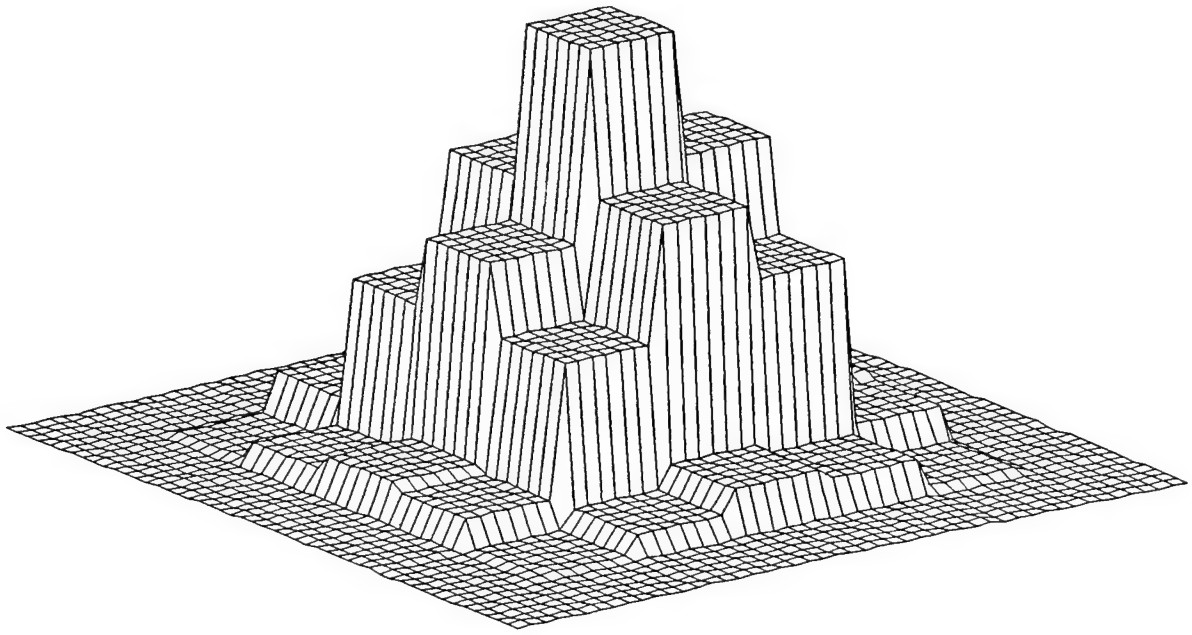

Fig. 4. Image at 0.8 waves defocus

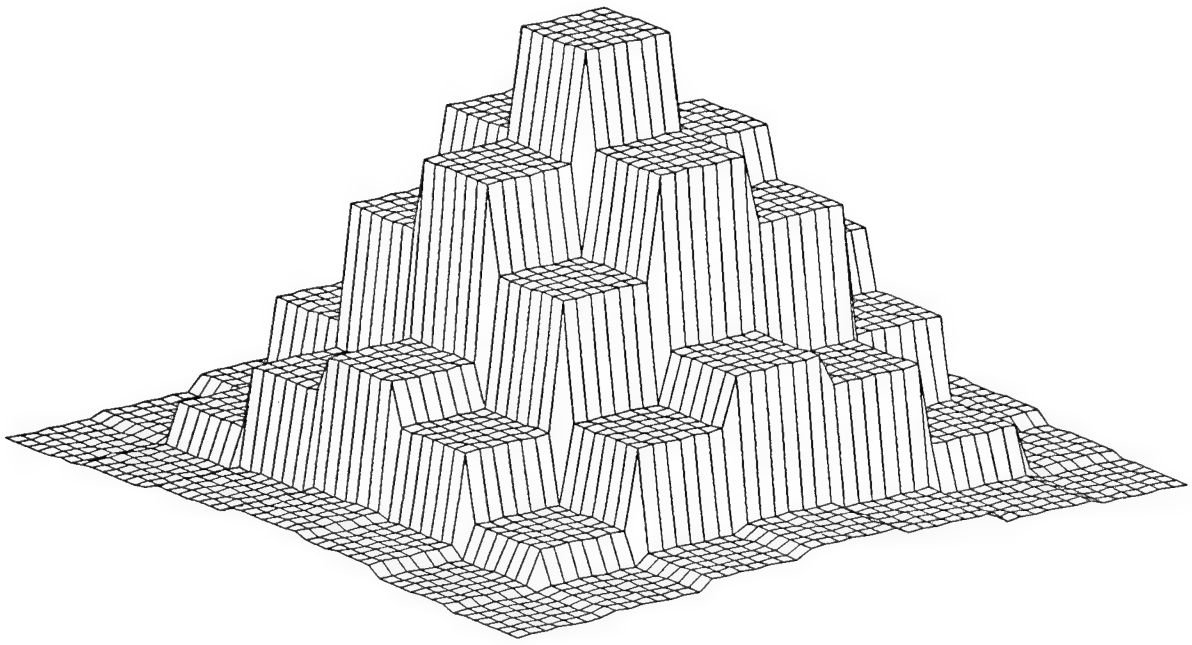

Fig. 5. Image at 1.6 waves defocus

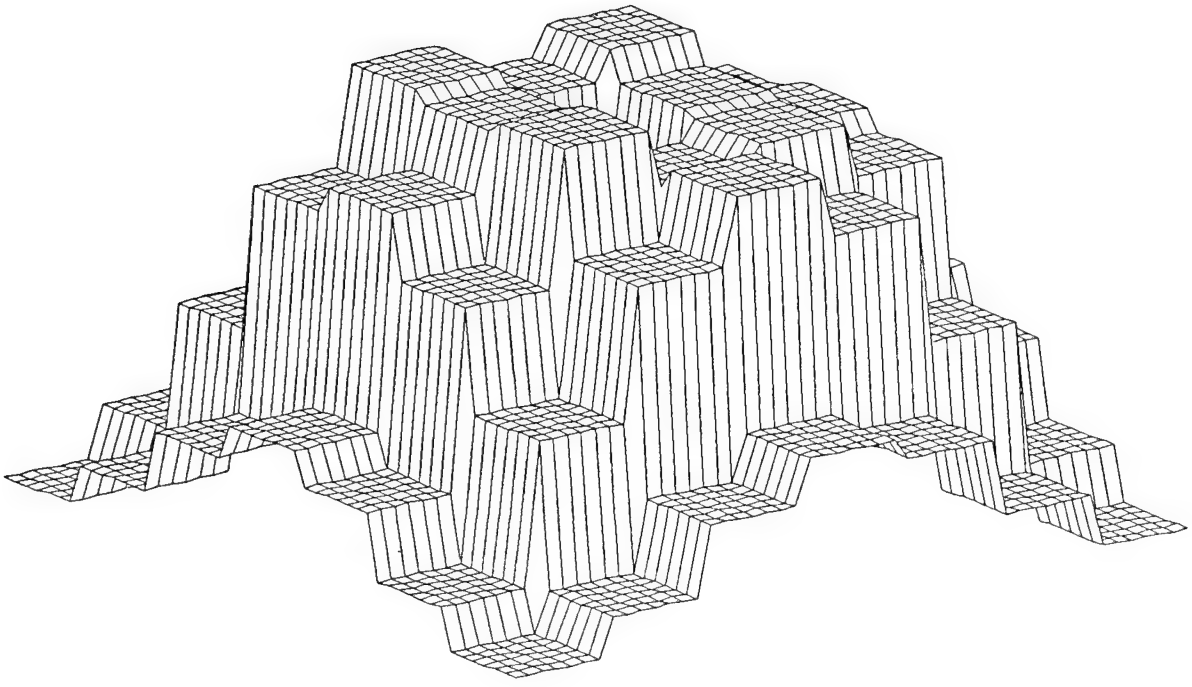

Fig. 6. Image at 2.4 waves defocus

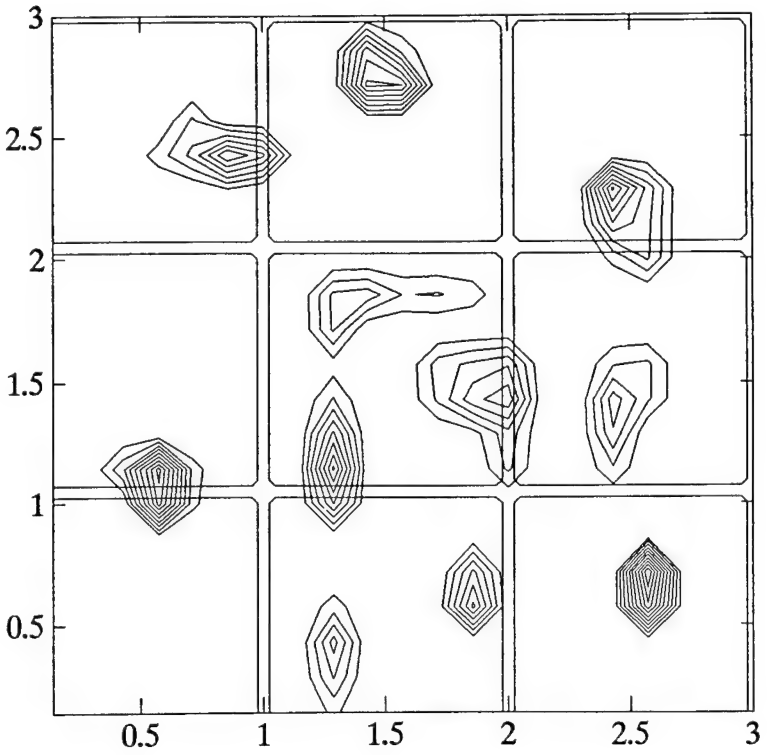

Fig. 7. Reconstruction from image with 1% added noise (beta = 5×10^{-5})

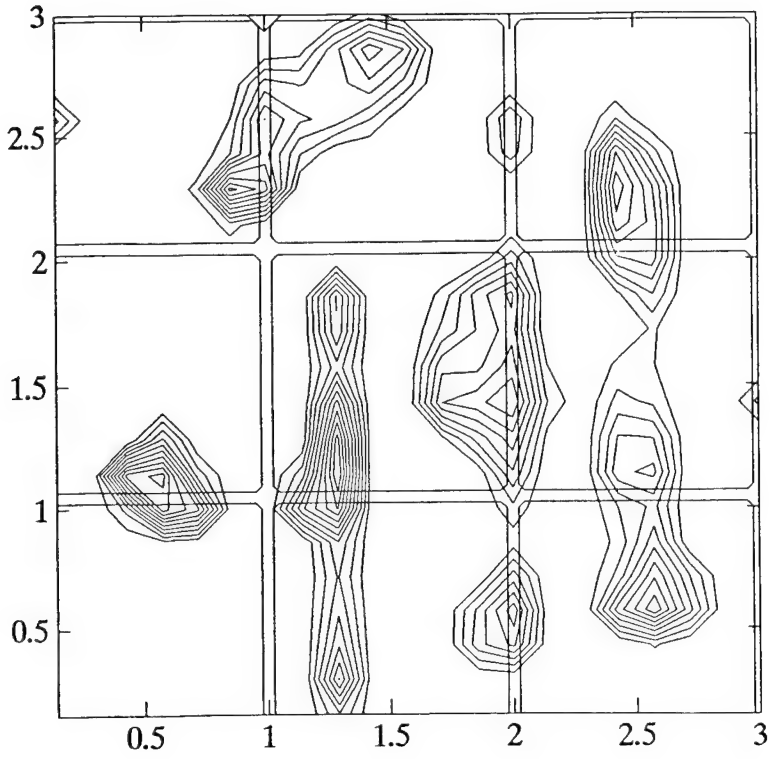

Fig. 8. Reconstruction from image with 5% added noise (beta = 10^{-3})

PROCEEDINGS
SPIE—The International Society for Optical Engineering

Inverse Problems in Scattering and Imaging

Michael A. Fiddy
Chair/Editor

20–22 July 1992
San Diego, California

Sponsored and Published by
SPIE—The International Society for Optical Engineering

Cooperating Organization
Society for Industrial and Applied Mathematics

Volume 1767

FAST IMAGE RECONSTRUCTION BASED ON THE REGULARIZED PSEUDOINVERSE OF THE IMAGING OPERATOR

John B. Abbiss

TITAN Spectron
1582 Parkway Loop
Suite B
Tustin, CA 92680-6505

Jeffery C. Allen, Richard P. Bocker, Harper J. Whitehouse

Naval Command, Control and Ocean Surveillance Center
RDT&E Division
San Diego, CA 92152-5000

ABSTRACT

A new algorithm for the restoration of extended images, the Regularized Pseudoinverse Deconvolution (RAPID) algorithm, is proposed. The algorithm consists of expanding the regularized pseudoinverse of the imaging operator into a sequence of terms which can be easily implemented using Fourier processing techniques. The first term of the expansion is closely related to generalized Wiener filtering if the point spread function is shift-invariant. The other terms in the expansion are correction terms which are small when the point spread function is shift-invariant, as is the case with many imaging systems. Even when the point spread function of the imaging system is space-variant, such as with a partially obscured imaging system or a system with severe aberrations, the correction terms are both few in number and easily implemented.

1. INTRODUCTION

In the absence of any other degrading effects, the performance of an optical system is ultimately restricted by diffraction. The finite extent of the entrance pupil imposes a fundamental upper limit on the system's spatial frequency response. The image quality of most operational systems will not, however, approach this theoretical limit very closely. It is possible that the design or construction will be flawed, as in the case of the Hubble Space Telescope, through defective manufacture, assembly or quality assurance procedures. In addition, aging of components will almost certainly compromise sensor performance at some stage in its lifetime, and, as in spaceborne operation, replacement may not be a simple task. The detector itself may impose limitations; for example, where a CCD array is used, information is lost in the inter-pixel areas, and image energy is integrated over the active area of each pixel. Other degrading factors will include defective pixels, noise in the CCD array and electronic subsystems, and a possibly obtrusive background. The methods of image restoration considered here were originally aimed at achieving performance beyond the diffraction limit [1], but are in fact capable of compensating simultaneously or separately for aberrations induced by the optical components and for the limitations of the detector. They are inherently robust and possess valuable noise-suppressing properties.

0-8194-0940-5/92/$4.00

The assumption is made that the overall effect of the optics can be described as a possibly time- and shift-variant blurring of the image due to diffraction and aberrations. Thus, at any instant of time, the point spread function may change across the sensor field-of-view. It can be assumed, however, that at any given point in the image the point spread function is effectively determined by its spatial location and is time-invariant over the integration time of the sensor and then undergoes a change at a later time. This is the case with the Hubble Space Telescope. It will be assumed that the set of point spread functions is known or can be measured. For the Hubble wide-field planetary camera there are two primary components to the point spread function; one due to diffraction by the aperture obstructions and the other due to spherical aberration. The point spread function is observed to be locally shift-invariant, and the image can be considered to be created by the summation across the entire field-of-view of the segmented set of localized point spread functions convolved with objects in the corresponding parts of the field.

However, for a simple refracting telescope designed for observation of extended objects, the aberrations causing image degradation can be highly space-variant. For an extended object, image segmentation introduces undesirable edge effects at the block boundaries when locally shift-invariant approximations of the globally shift-variant point spread function are used to process the separate blocks. In addition, post processing interpolation or iteration is required to smooth the block boundaries in the final composite image [2]. In the approach proposed in this paper, the point spread function is allowed to vary continuously over the whole image and a special decomposition of the image reconstruction operator is performed which permits Fourier techniques to be used to process the entire image.

The overall optical system imaging equation can be written as a Fredholm integral equation of the first kind. Solutions to ill-posed problems of this type are known to be numerically unstable [3, 4]. Additionally, it is anticipated that the image will be spatially sampled by a solid-state sensor which will introduce spatial integration, discretization and associated noise processes. Thus, after scanning the image into a vector, the integral representing the continuous image can be rewritten as a matrix expression. In general, the presence of the sensor noise takes the measured image vector out of the span of the columns of the kernel matrix, which is typically highly ill-conditioned. Thus, when it is desired to estimate what the image would have been in the absence of aberrations or with less diffraction than the instrument can provide, techniques derived from regularization theory are required to restore stability to the reconstruction. By introducing a suitable error criterion (based on, *e.g.*, a vector norm), images can be constructed which are, in terms of the chosen criterion, closer to the undistorted geometrical image of the object than the detected image data.

To the extent permitted by the noise in the image, in-band effects can usually be removed by some form of pseudoinverse filter [5]. However, detector pixellation and the finite aperture of any system set resolution limits not so easily overcome, and a method for achieving spectral extrapolation has to be devised. The spatial spectrum of the object is the Fourier transform of its amplitude, in the coherent case, or its intensity, in the incoherent case. If the object is known to be of finite extent, its Fourier transform is an analytic function, and the out-of-band part of the spectrum can in principle be fully recovered by analytic continuation [6] of the image spectrum after removal of any in-band distortion. The inverse Fourier transform of this extended spectrum would then yield a perfect image of the original object. Equivalently, one could attempt to solve directly the equation describing the imaging process. This, however, involves the inversion of an ill-conditioned matrix, and the restoration process is intrinsically unstable, even small amounts of noise rendering the results meaningless. These difficulties may be surmounted by applying the methods of regularization theory [7], developed to deal with ill-posed problems of this type; the solution is derived by means of a constrained least-squares procedure in which a regularization parameter plays an essential role. Stability in the restored image, which is computed via the regularized pseudoinverse of the imaging matrix, is controlled by this parameter. Its optimal value depends on the signal-to-noise ratio in the data.

2. NATURE OF THE PROBLEM

We wish to estimate the object f from an image g, given that

$$g = Af + r \tag{1}$$

where A is the imaging operator and r represents the corrupting effect of additive noise. For clarity in the analysis, we consider the one-dimensional case with the operator A given in integral form by

$$(Af)(y) = \int_a^b A(x,y) f(x)\, dx, \qquad c \le y \le d. \tag{2}$$

In the absence of noise, Eq.(1) becomes a Fredholm equation of the first kind, in which the unknown function appears only under the integral sign. We can identify the sources of difficulty in solving this equation in the presence of noise or other perturbations (such as computer round-off error) by means of a singular function analysis [8].

We expand the kernel of the integral in terms of the singular functions $u_i(x)$ and $v_i(y)$, orthonormal systems in object and image spaces respectively, and the singular values σ_i:

$$A(x,y) = \sum_{i=1}^{\infty} \sigma_i\, u_i(x)\, v_i(y). \tag{3}$$

The object and image can be expanded in the forms

$$f(x) = \sum_{i=1}^{\infty} f_i\, u_i(x) \tag{4}$$

$$g(y) = \sum_{i=1}^{\infty} g_i\, v_i(y) \tag{5}$$

where the coefficients are related to $f(x)$ and $g(y)$ by the integral formulae

$$f_i = \int_a^b f(x)\, u_i(x)\, dx \tag{6}$$

and

$$g_i = \int_c^d g(y)\, v_i(y)\, dy. \tag{7}$$

In the noiseless case $(r = 0)$, we find from Eqs. (2) and (6)

$$(Af)(y) = \sum_{i=1}^{\infty} \sigma_i\, f_i\, v_i(y). \tag{8}$$

It follows, using Eqs.(5) and (8), that

$$g_i = \sigma_i f_i \tag{9}$$

and hence

$$f(x) = \sum_{i=1}^{\infty} \frac{g_i}{\sigma_i} u_i(x). \tag{10}$$

Thus the object function can in principle be perfectly reconstructed from the set $\{g_i\}$ of image coefficients.

Now consider the effects of noise. By expanding $r(y)$ in terms of the $v_i(y)$, we can derive the contribution of the noise to the new image coefficients:

$$g'_i = \sigma_i f_i + r_i. \tag{11}$$

The estimate of $\hat{f}(x)$ is now

$$\hat{f}(x) = f(x) + \sum_{i=1}^{\infty} \frac{r_i}{\sigma_i} u_i(x). \tag{12}$$

Image formation is characteristically described by an integral transform of convolution type, *i.e.*, A is a convolution operator. Its singular-value spectrum typically decays asymptotically at an exponential rate[9]. Since the r_i will in general decrease less quickly, the sum in Eq.(12) will be divergent and no bound will exist for the 'distance' (in the sense of some appropriate metric) between the true object and the reconstruction. The effect of the noise on the reconstructed image is a manifestation of the fact that convolution is a strongly smoothing process - closely similar images can correspond to widely differing objects. Thus image restoration is an ill-posed problem, small perturbations in the data causing large changes in the solution represented by Eq.(12).

3. THE REGULARIZED SOLUTION

The methods of regularization theory[7, 10] can be exploited to convert the problem to a related well-posed one, *i.e.*, one for which the solution exists, is unique and depends continuously on the data. Since we shall later be concerned with computed reconstructions, we henceforth consider the problem in its finite discretized form; thus the imaging operator A becomes a matrix. Although strictly we should introduce new symbols, for convenience we continue to use f, g, and A to denote the discrete forms of object, image, and imaging operator. Generally $f \in C^n$, $g \in C^m$ and $A \in C^{m \times n}$.

To regularize the problem, we shall modify A. We impose constraints[11] on possible solutions f' by requiring that

$$\| Af' - g \|^2 \leq \varepsilon^2 \tag{13}$$

where ε is some suitable measure of the noise in the image, and that

$$\| f' \|^2 \leq \xi^2 \tag{14}$$

where ξ is some suitable measure of the permitted 'signal strength' of the solution. ($\| \bullet \|$ denotes norm in the Hilbert spaces associated with object and image.) We combine these constraints and minimize

$$\| Af' - g \|^2 + \beta \| f' \|^2$$

where the regularization parameter β is given by

$$\beta = \varepsilon^2 / \xi^2 \quad . \tag{15}$$

The minimum-norm solution to this constrained least-squares problem is given by

$$f_\beta = A_\beta^+ g \tag{16}$$

where

$$A_\beta^+ = (A^H A + \beta I)^{-1} A^H. \tag{17}$$

We note the relationship of A_β^+ (which we shall call the regularized pseudoinverse) to A^+, the Moore-Penrose pseudoinverse [12]

$$A^+ = \lim_{\beta \to 0} A_\beta^+. \tag{18}$$

The inverse of $(A^H A + \beta I)$ always exists, since $A^H A$ is non-negative definite and the regularization parameter β is positive. A value of β should be chosen which balances data fidelity against smoothness in the reconstruction. Methods are also available for determining β from the image data themselves [13, 14].

4. SINGULAR VALUE AND FOURIER DECOMPOSITIONS

It is often convenient to compute the regularized pseudoinverse via the singular value decomposition (SVD) of A:

$$A = U \Sigma V^H \tag{19}$$

where [15]

$$U^H U = I_{m'} \qquad V^H V = V V^H = I_n \tag{20}$$

and

$$\Sigma = diag\,(\sigma_1, \sigma_2, \ldots, \sigma_n), \qquad \sigma_i \geq 0. \tag{21}$$

The singular values $\{\sigma_i\}$ are assumed to have been arranged in descending order of magnitude

$$\sigma_1 > \sigma_2 > \sigma_3 > \ldots > \sigma_n. \tag{22}$$

Then we find

$$f_\beta = V \; \Sigma_\beta^+ \; U^H \; g \tag{23}$$

where

$$\Sigma_\beta^+ = diag\left(\ldots, \frac{\sigma_i}{\sigma_i^2 + \beta}, \ldots \right). \tag{24}$$

This representation is useful when the behavior of the reconstruction as a function of the regularization parameter is being studied. Regularization can equivalently be achieved by setting β to zero, and simply truncating the singular value series at a point which is dependent on the noise level [16].

An *ab initio* computation of f_β via Eq.(23) requires the SVD of A followed by two matrix-vector multiplications. If the regularized pseudoinverse can be precomputed, only one matrix-vector product is needed to generate the reconstruction. For images of more than modest sizes, however, the computation rapidly becomes burdensome. If, for example, f and g are 100x100, A is a 10^4-by-10^4 matrix, and the matrix-vector product requires 10^8 multiplications. There will also be considerable storage demands. Thus, if major computational resources are not available, some means of simplifying the calculation will be needed in many cases of practical interest.

If the matrix A were square circulant of order n ($A = [a_{j-i+1}]$, subscript mod n), we could dramatically reduce the computational burden by exploiting the fact that the Fourier transform diagonalizes a circulant [17]. If F denotes the Fourier matrix:

$$F^H = \frac{1}{\sqrt{n}} \begin{bmatrix} 1 & 1 & 1 & \ldots & 1 \\ 1 & w & w^2 & \ldots & w^{n-1} \\ 1 & w^2 & w^4 & \ldots & w^{2(n-1)} \\ & & & \ldots & \\ 1 & w^{n-1} & w^{2(n-1)} & \ldots & w^{(n-1)^2} \end{bmatrix}, \qquad w = exp(i2\pi/n) \tag{25}$$

then

$$A = F^H \Lambda F \tag{26}$$

where

$$\Lambda = diag\,(\lambda_1, \lambda_2, \ldots, \lambda_n). \tag{27}$$

It follows from Eq.(16) that

$$f_\beta = F^H \; \Lambda_\beta^+ \; F \; g \tag{28}$$

where

$$\Lambda_\beta^+ = diag\left(\ldots, \frac{\bar{\lambda}_i}{|\lambda_i|^2 + \beta}, \ldots \right) \tag{29}$$

and $\bar{\lambda}_i$ is the complex conjugate of λ_i. We note that the singular values $\{\sigma_i\}$ and the eigenvalues $\{\lambda_i\}$ are related by

$$\sigma_i = |\lambda_i|. \tag{30}$$

The operational counts for the SVD and the FFT are $O(n^3)$ and $O(n \log n)$, respectively.

5. STRUCTURE OF THE IMAGING MATRIX

Under some circumstances the imaging matrix can be readily modified to circulant form. For a shift-invariant one-dimensional system, the image is a convolution of the point spread function and the object. If the sampling intervals in image and reconstruction spaces are equal, the matrix A is then Toeplitz ($A = [a_{j-i}]$). If, in addition, the image and reconstruction vectors are of equal length, A can be padded to circulant form. In two dimensions, f and g are matrices and must be mapped into vectors. The manner in which this is performed will determine the structure of A. For column-wise mapping, for instance, again with equal sampling in image and reconstruction spaces, A becomes block-Toeplitz with Toeplitz blocks. If the image and reconstruction matrices have the same number of elements, A can be padded to become block-circulant with circulant blocks, and the problem is again amenable to Fourier transform methods. In both one and two dimensions, it should be noted that the penalty associated with the expansion of A to circulant or block-circulant form is a relaxation of the support constraint on the reconstruction, which renders the calculation, and in particular the degree of resolution enhancement achieved, much more sensitive to noise [18]. An alternative construction in the two-dimensional case is to zero-pad f and g into larger matrices and then to use a linear congruential scan to map the padded matrices into vectors. A then becomes a circulant matrix since the linear congruential scan is an isomorphism between 2D convolution and 1D convolution [19].

For less structured imaging matrices (*e.g.*, if the system is weakly space-variant) it may be asked whether accelerated computation of the matrix-vector product is still possible. In this context, recent work on circulant approximations to matrices of quite general form [20] appears highly relevant, and includes the following result. For any matrix, A_β^+ say, we can write

$$A_\beta^+ = C_o + \sum_{m=1}^{\alpha} L(x_m)\, C^T(y_m) \tag{31}$$

where C_o is a circulant matrix with the same last row as A_β^+, $L(x_m)$ is a lower triangular Toeplitz matrix with x_m as its first column, and $C(y_m)$ is a circulant matrix whose last row is y_m. The $\{x_m\}$ and $\{y_m\}$ may be obtained from the truncated SVD of the cyclic displacement of A_β^+:

$$A_\beta^+ - EA_\beta^+E^T = \sum_{m=1}^{\alpha} x_m y_m^T \tag{32}$$

where E is the cyclic downshift matrix

$$E = \begin{bmatrix} 0\,0\,0 \ldots 0\,1 \\ 1\,0\,0 \ldots 0\,0 \\ 0\,1\,0 \ldots 0\,0 \\ \quad\quad \ldots \\ 0\,0\,0 \ldots 1\,0 \end{bmatrix}. \tag{33}$$

For imaging matrices with strongly Toeplitz features, α should be a small number.

6. THE REGULARIZED PSEUDOINVERSE DECONVOLUTION ALGORITHM

Consider a two-dimensional optical imaging system whose point spread function is both time- and space-invariant. In the presence of additive noise, the imaging equation connecting the input (extended object), the output (degraded image), and the point spread function (impulse response) is given by the following two-dimensional convolutional integral equation

$$i(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} o(x',y')\,p(x-x',y-y')\,dx'dy' + n(x,y)\,. \tag{34}$$

In shorthand notation

$$i = p\bigstar\bigstar o + n \tag{35}$$

where $o(x,y)$ represents the extended object, $i(x,y)$ the degraded image, $p(x,y)$ the point spread function, and $n(x,y)$ the noise. It is assumed that the noise is independent of position in the image.

The discrete version of Eq.(34) can, of course, be cast[21] into the following vector-matrix form

$$g = Af + r. \tag{36}$$

When, in particular, f, g, and r represent the one-dimensional column vectors formed by stacking the rows or columns of the discretized versions of the input $o(x,y)$, output $i(x,y)$, and the noise $n(x,y)$, respectively, the two-dimensional imaging matrix A is block-Toeplitz with Toeplitz blocks. It can be shown, using known results [22] that the regularized pseudoinverse A_β^+ is block-persymmetric with persymmetric blocks. The inverse of a Toeplitz matrix is persymmetric, and persymmetric matrices obtained by inverting Toeplitz matrices have much more Toeplitz-like structure than general persymmetric matrices. In particular, their displacement rank is the same as that of the parent-Toeplitz matrix[23, 24]. Displacement rank, it should be noted, is a quantitative measure of the closeness of a given matrix to being Toeplitz. In one dimension, the displacement rank of a Toeplitz matrix is ≤ 2, and the displacement rank of A_β^+ is ≤ 4.

In the early stages of this investigation, one of the authors (R.P.B.) noted that for a variety of point spread functions being studied the corresponding computer-generated regularized pseudoinverse matrices appeared to have banded block-Toeplitz structure. He considered the implication of this observation. In particular, if a space-invariant point spread function used in a two-dimensional linear convolutional imaging equation leads to a block-Toeplitz imaging matrix with Toeplitz blocks, then the converse must also be true. That is, given a block-Toeplitz imaging matrix containing Toeplitz blocks, then the corresponding space-invariant point spread function which gave rise to this imaging matrix could be easily ascertained. This implies that from the regularized pseudoinverse an inverse point spread function $d_\beta(x,y)$ could be constructed which could be used to process the image $i(x,y)$ and form an estimate $\hat{o}_\beta(x,y)$ of the original object. This two-dimensional linear convolution technique is summarized by the equation

$$\hat{o}_\beta = d_\beta\bigstar\bigstar i. \tag{37}$$

The technique was tested on a digital computer with encouraging results. Further experimental investigations indicate that results obtained with this Regularized Pseudoinverse Deconvolution (RAPID) algorithm are comparable in quality to those obtained using parametric Wiener filtering. An error analysis is given in Appendix A. Results using degraded images processed with both the RAPID algorithm and parametric Wiener filtering are presented in section 8.

7. DECONVOLUTION VIA WIENER FILTERING

Wiener filtering [21] is a well-known technique for processing images degraded and corrupted by noise as described by Eq.(34). From a knowledge of the point spread function $p(x,y)$ characterizing the optical imaging system, it is possible to compute the corresponding optical transfer function $\tilde{p}(f_x, f_y)$ using the two-dimensional Fourier transform

$$\tilde{p}(f_x, f_y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(x,y)\, exp\,[-i2\pi(xf_x + yf_y)]\, dxdy. \tag{38}$$

From the optical transfer function and knowledge of the power spectra of object $S_o(f_x, f_y)$ and noise $S_n(f_x, f_y)$, the parametric Wiener filter can be constructed, namely

$$\tilde{w}_\gamma(f_x, f_y) = \frac{\left|\tilde{p}(f_x, f_y)\right|^2}{\tilde{p}(f_x, f_y)\left[\left|\tilde{p}(f_x, f_y)\right|^2 + \gamma\left|\frac{S_n(f_x, f_y)}{S_o(f_x, f_y)}\right|\right]}. \tag{39}$$

When $\gamma = 1$, Eq.(39) reduces simply to the Wiener filter. If γ is variable we refer to this as the parametric Wiener filter. In the absence of noise, either form of the Wiener filter reduces to the ideal inverse filter. When the power spectra are not known, which is often the case in practice, Eq.(39) can be approximated by

$$\tilde{w}_\gamma(f_x, f_y) \approx \frac{\left|\tilde{p}(f_x, f_y)\right|^2}{\tilde{p}(f_x, f_y)\left[\left|\tilde{p}(f_x, f_y)\right|^2 + \gamma\right]}. \tag{40}$$

From the Wiener filter, using either Eqs.(39) or (40), an inverse point spread function $w_\gamma(x,y)$ can be constructed using the two-dimensional inverse Fourier transform. That is,

$$w_\gamma(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \tilde{w}_\gamma(f_x, f_y)\, exp\,[+i2\pi(xf_x + yf_y)]\, df_x df_y. \tag{41}$$

With the inverse point spread function given by Eq.(41), an estimate $\hat{o}_\gamma(x,y)$ of the object can be computed using the two-dimensional linear convolutional equation

$$\hat{o}_\gamma = w_\gamma \star\star i \tag{42}$$

where, again, $i(x,y)$ is the degraded image.

8. RESULTS

Preliminary results obtained using the RAPID algorithm are presented in this section. For purposes of comparison, results obtained using the parametric Wiener filter algorithm are also included. The optical system considered for these studies was a simple spherical converging lens as shown in Fig. 1.

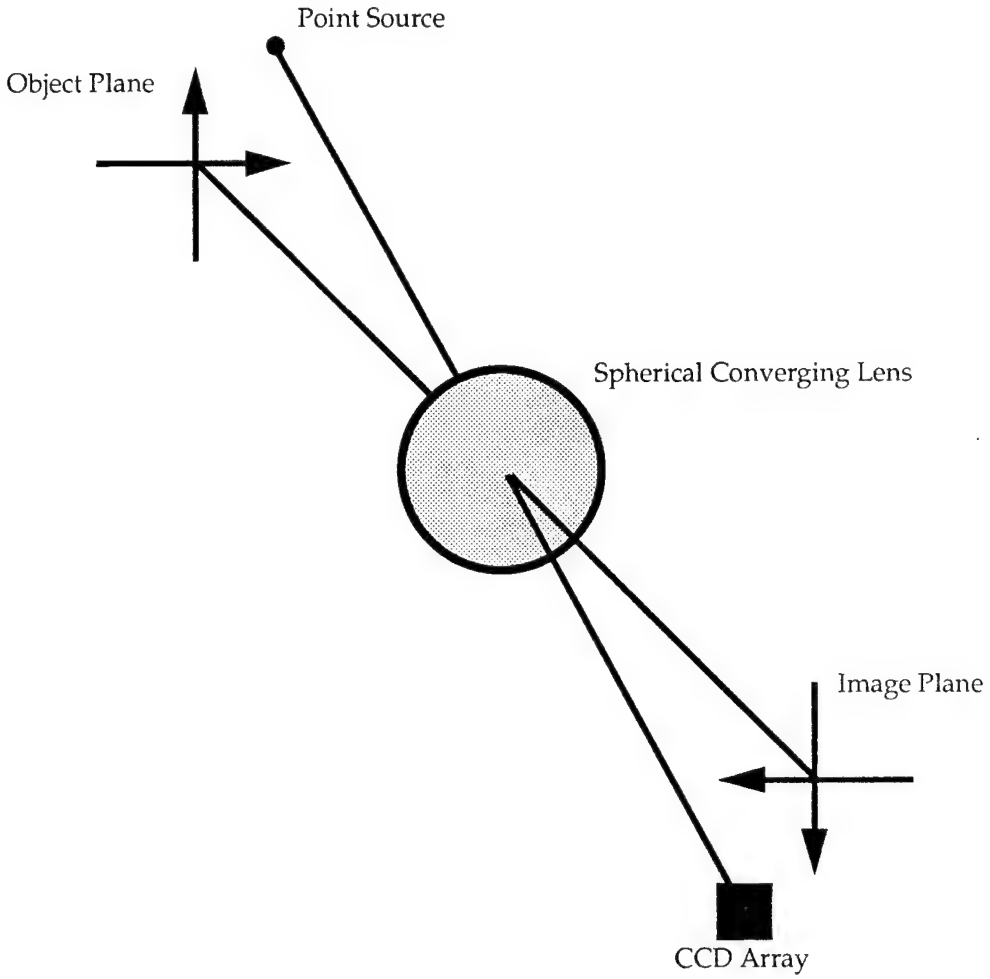

Figure 1. A Simple Spherical Converging Lens Imaging System

The object and image planes are coplanar and orthogonal to the optical axis of the lens. An arbitrary point source located in the object plane gives rise to an intensity distribution (the point spread function) in the image plane. A charge-coupled device (CCD), for example, can be used to measure the point spread function. A ray-trace program was used in the synthesis of four different point spread functions dominated by spherical aberration (Fig. 2), coma (Fig. 3), astigmatism (Fig. 4), and defocus (Fig. 5). The object plane distance (mm), image plane distance (mm), focal length (mm), F-number, tangential field-angle (deg.), and sagittal field-angle (deg.) associated with each of these four figures are summarized in Table 1. The same CCD model was used in all simulations. The array consisted of a 31-by-31 planar arrangement of square detectors measuring 0.01 mm on a side with a center-to-center spacing of 0.01 mm. Each point spread function was obtained by tracing 20,000 rays.

Table 1

	Spherical Aberration	Coma	Astigmatism	Defocus
Object plane distance	48.6	48.6	48.6	48.6
Image plane distance	52.0	52.0	50.4	51.0
Focal length	24.3	24.3	24.3	24.3
F-number	4.80	5.70	5.70	4.00
Tangential field-angle	0.00	2.87	5.91	0.00
Sagittal field-angle	0.00	2.87	0.00	0.00

On the top line, center diagram, of Figs. 2 through 5 are mesh plots of the four point spread functions considered. Each point spread function $p(x,y)$ is represented by a 31-by-31 matrix. The ideal extended object $o(x,y)$ used in this analysis was a 256-by-256 matrix which is displayed as an 8-bit gray-level diagram in the upper-left hand corner of each of these figures. Performing a two-dimensional convolution, see Eq. (35), of the extended object with the point spread function (in the absence of noise) yields a 286-by-286 degraded image $i(x,y)$ of which the central 256-by-256 portion of the degraded image is shown in the upper-right hand corner of each of these figures. The full 286-by-286 matrix is used, however, in subsequent image processing computations.

On line two of each of the four figures are mesh plots of the inverse point spread functions (31-by-31 matrices) $d_\beta(x,y)$ (left-diagram) and $w_\gamma(x,y)$ (right-diagram) obtained using the RAPID and Wiener filter algorithms, respectively. Performing a two-dimensional convolution of the point spread function $p(x,y)$ with each of the inverse point spread functions $d_\beta(x,y)$ and $w_\gamma(x,y)$ yields processed point spread functions (61-by-61 matrices). The central 31-by-31 portions of these processed point spread functions are shown as mesh plots on line three (left- and right-diagrams, respectively). For purposes of comparison, a 31-by-31 null-array with a single non-zero entry for the center pixel (Delta) is also displayed on line three, center diagram.

The values of the regularization parameters (β and γ) which gave rise to the best processed point spread functions, judged visually, for this analysis are: spherical aberration (0 and 0), coma ($3x10^{-2}$ and $1x10^{-4}$), astigmatism ($2x10^{-3}$ and $2x10^{-2}$), and defocus ($5x10^{-4}$ and $1x10^{-2}$). These same regularization parameter values were used in computing the object estimates $\hat{o}_\beta(x,y)$ and $\hat{o}_\gamma(x,y)$ using Eqs. (37) and (42), respectively. These object estimates (316-by-316 processed images) are displayed as gray-level diagrams on the bottom-line of each of the four figures. Only the central 256-by-256 portions of the processed images are shown.

The results presented in Figs. 2 through 5 were based on studies using synthesized point spread functions. We were fortunate to obtain real digitized degraded images of the planet Saturn taken with the wide-field planetary camera of the Hubble Space Telescope. The upper diagram in Fig. 6 shows a 400-by-250 degraded (unprocessed) image of the planet Saturn. The second line in Fig. 6 shows two 31-by-31 degraded images of different stars (called star #1 and star #2) also taken with the wide-field planetary camera. The RAPID algorithm was used to process the degraded image of Saturn using the two star images as point spread functions characterizing the degradation process. In particular, star #1 image was used as the input point spread function. Both star #1 and star #2 images were first processed using the inverse point spread function obtained using the star #1 image only. The regularization parameter, β, selected was the one which gave rise to processed star images of equal quality, based on a minimum entropy criterion. This same inverse point spread function was then used, via Eq.(35), to process the degraded Saturn image. The size of the reconstructed image was 430-by-280. The central 400-by-250 portion of this reconstruction is shown in the lower diagram of Fig. 6.

Ideal

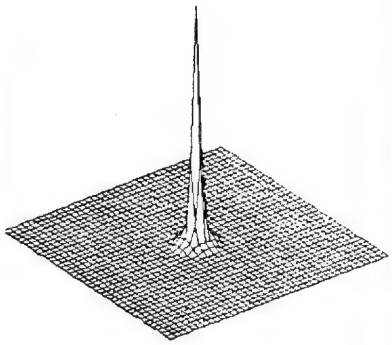

Point Spread Function

Degraded

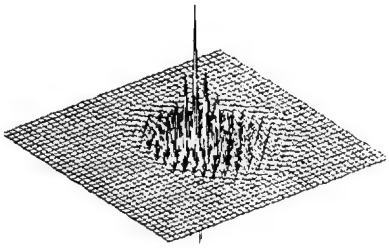

Inverse PSF - Pseudoinverse

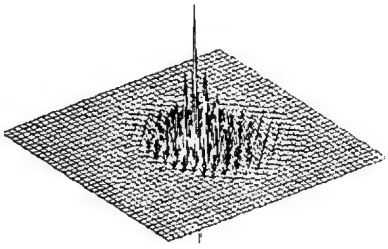

Inverse PSF - Wiener

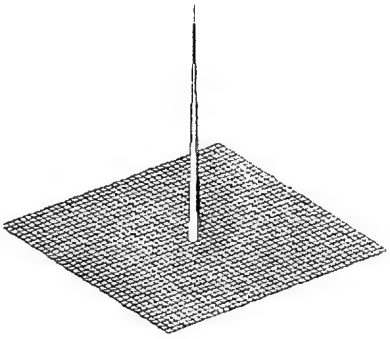

Processed PSF - Pseudoinverse

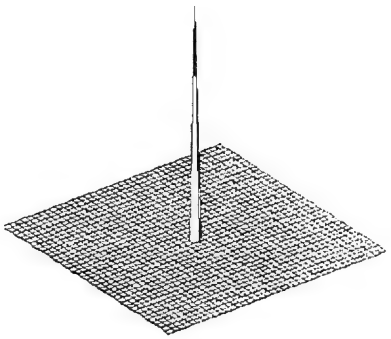

Delta

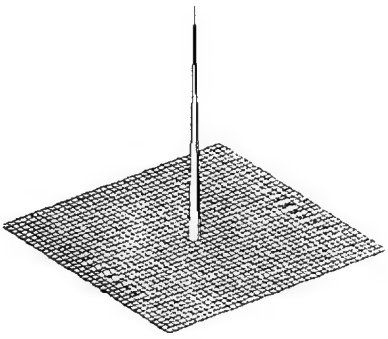

Processed PSF - Wiener

Processed - Pseudoinverse

Processed - Wiener

Figure 2. Point Spread Function Dominated by Spherical Aberration

Ideal

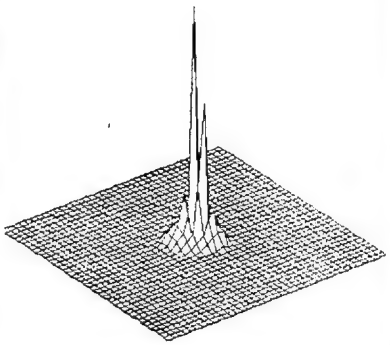

Point Spread Function

Degraded

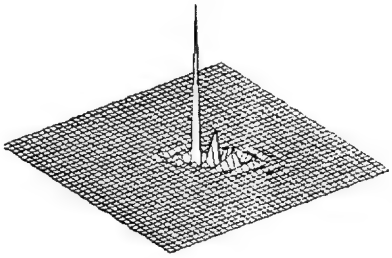

Inverse PSF - Pseudoinverse

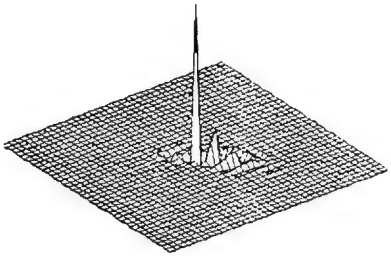

Inverse PSF - Wiener

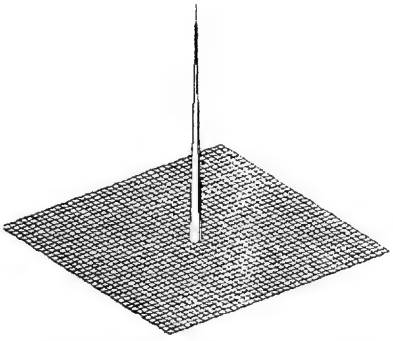

Processed PSF - Pseudoinverse

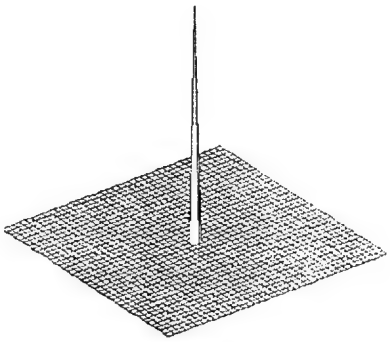

Delta

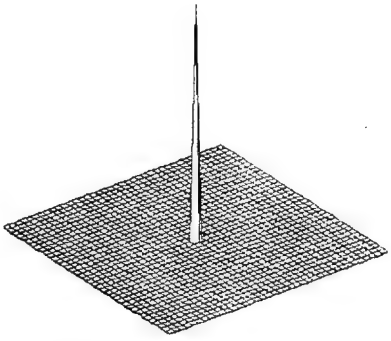

Processed PSF - Wiener

Processed - Pseudoinverse

Processed - Wiener

Figure 3. Point Spread Function Dominated by Coma

Ideal

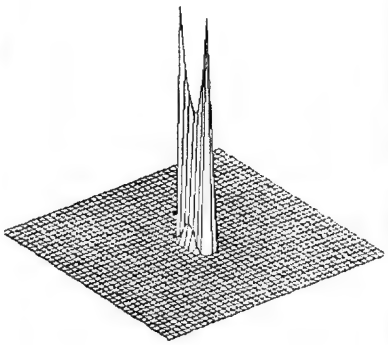

Point Spread Function

Degraded

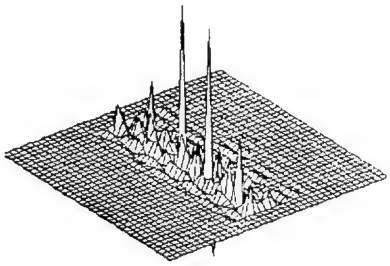

Inverse PSF - Pseudoinverse

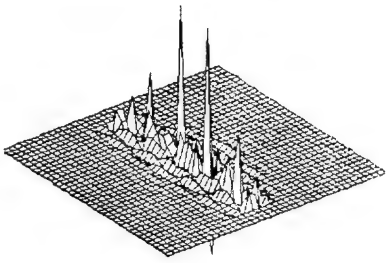

Inverse PSF - Wiener

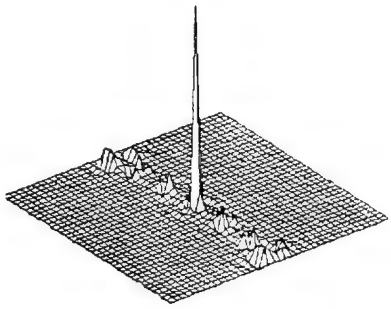

Processed PSF - Pseudoinverse

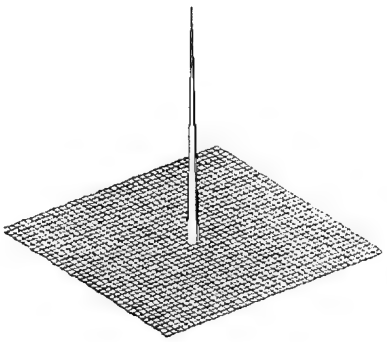

Delta

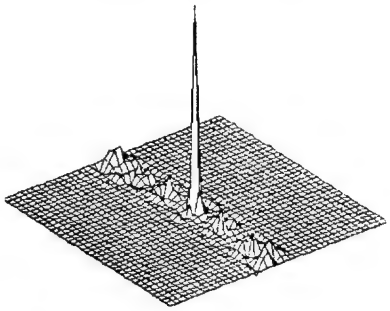

Processed PSF - Wiener

Processed - Pseudoinverse

Processed - Wiener

Figure 4. Point Spread Function Dominated by Astigmatism

Ideal

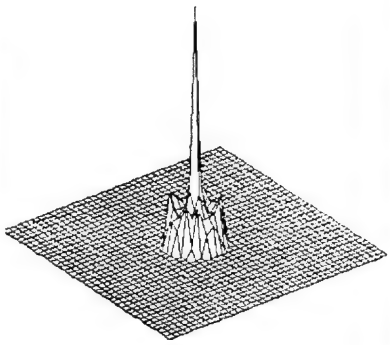

Point Spread Function

Degraded

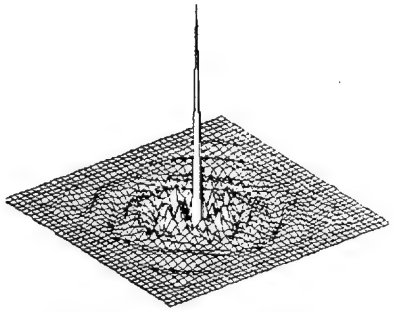

Inverse PSF - Pseudoinverse

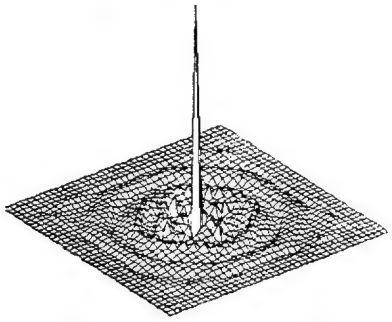

Inverse PSF - Wiener

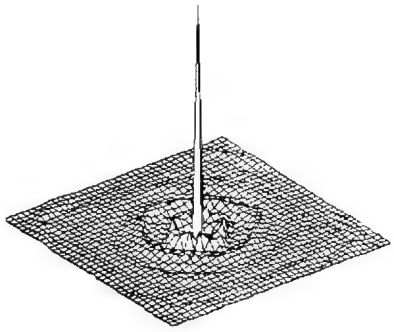

Processed PSF - Pseudoinverse

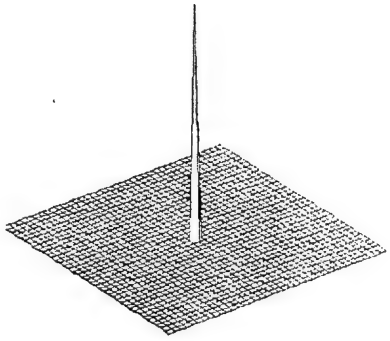

Delta

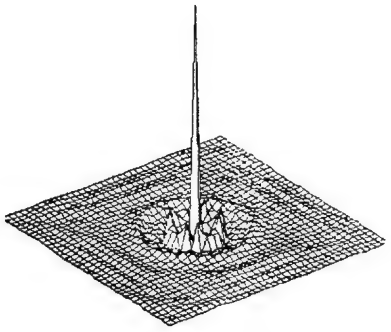

Processed PSF - Wiener

Processed - Pseudoinverse

Processed - Wiener

Figure 5. Point Spread Function Dominated by Defocus

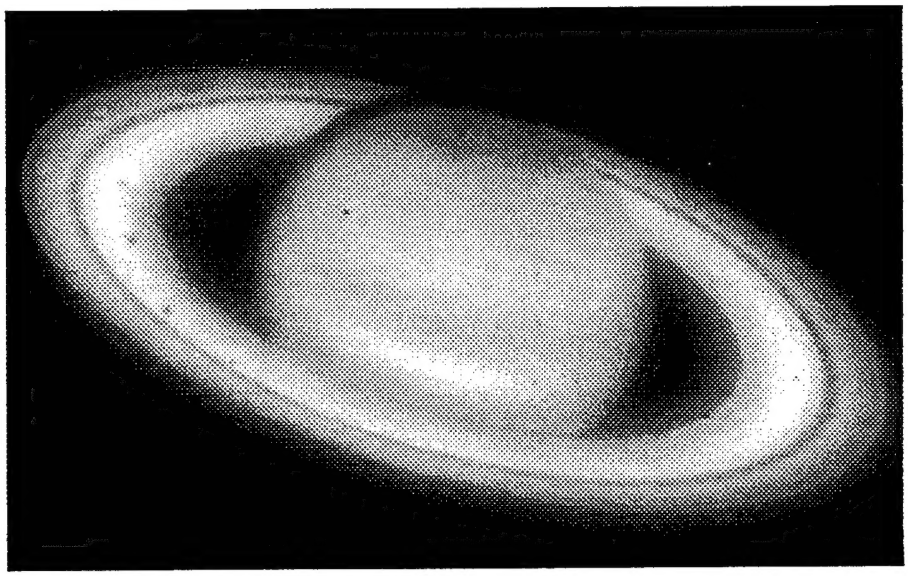

Unprocessed Saturn Image

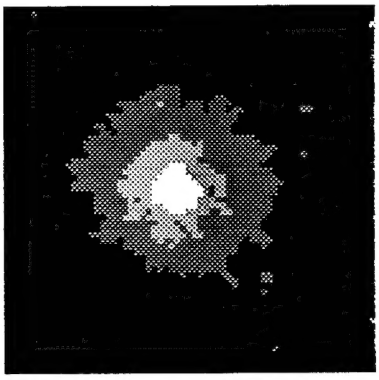

Star #1 Image

Star #2 Image

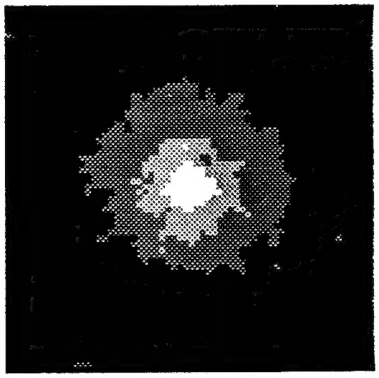

Processed Saturn Image

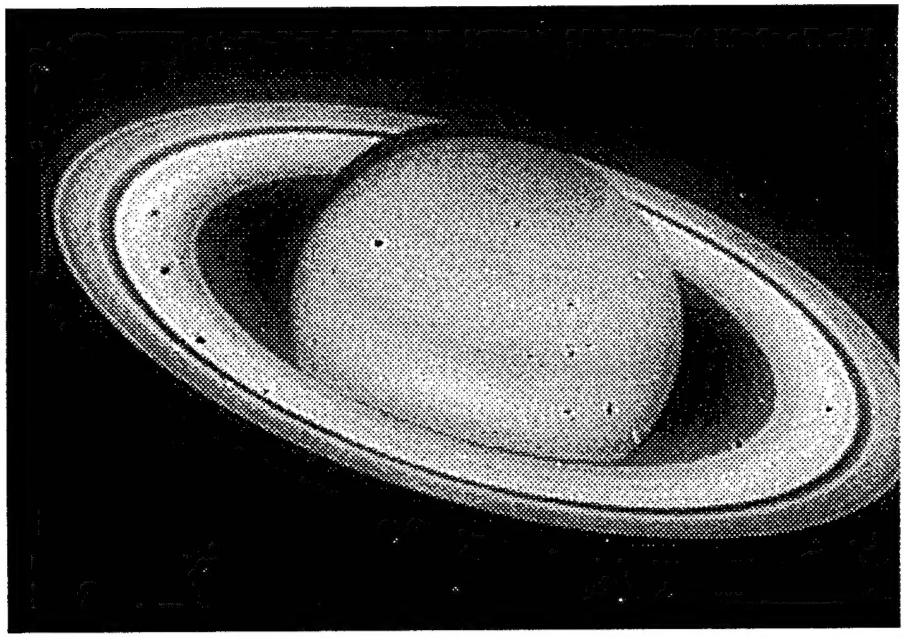

Figure 6. Hubble Space Telescope Imagery

9. CONCLUSIONS

A new algorithm, the Regularized Pseudoinverse Deconvolution (RAPID) algorithm, has been developed and has been shown to be equivalent in performance to Wiener filtering for shift-invariant point spread functions. The algorithm can be implemented either by direct convolution or indirectly using Fast Fourier Transform (FFT) techniques. The advantage of this approach is that, through the use of linear congruential scanning and the application of the circulant expansion of equation (31), regularized reconstruction methods can be applied to extended images degraded by shift-variant point spread functions.

10. ACKNOWLEDGEMENTS

This research was supported by the Strategic Defense Initiative Organization, Innovative Science and Technology Office, (SDIO/ISTO) through the Office of Naval Research and managed by Dr. Keith Bromley. Thanks are due to the National Aeronautics and Space Administration (NASA) for providing the digitized image of the planet Saturn obtained with the wide-field planetary camera of the Hubble Space Telescope. Special thanks go to Mr. Tom Phillips, Code 414, Naval Command, Control and Ocean Surveillance Center, RDT&E Division, for developing the ray-trace program and supplying the point spread functions used in the computations.

11. REFERENCES

1. J. B. Abbiss and P. G. Earwicker, "Compact operator equations, regularization and super-resolution," *Proceedings, IMA Conference on Mathematics in Signal Processing*, University of Bath, Bath, United Kingdom, 17-19 September 1985. (Published by Clarendon Press, Oxford, U.K., 1987.)
2. D. J. Lindler and S. R. Heap, "Block iterative restoration of astronomical images with the NASA/Goddard massively parallel processor," *Proceedings, Second International Conference on Supercomputing*, (Lana P. Kartashev and Steven I. Kartashev, eds.), International Supercomputing Institute, Inc., 1987.
3. P. A. Jansson (Ed.), *Deconvolution with Applications in Spectroscopy*, Academic Press, New York, 1984.
4. C. K. Rushforth, "Signal restoration, functional analysis, and Fredholm integral equations of the first kind," in: H. Stark (Ed.), *Image Recovery: Theory and Application*, Academic Press, New York, 1987.
5. A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, New Jersey, 1989.
6. M. A. Evgrafov, *Analytic Functions*, reprinted by Dover Publications, New York, 1978.
7. A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*, V. H. Winston & Sons, Washington DC, 1977.
8. F. Smithies, *Integral Equations*, Cambridge Tract No. 49, Cambridge University Press, Cambridge, England, 1958.
9. B. R. Frieden, "Evaluation, design and extrapolation methods for optical signals, based on the use of the prolate functions," *Progress in Optics*, (E. Wolf, ed.), Vol. IX, pp. 313-407, North Holland, Amsterdam, 1971.
10. C. W. Groetsch, *The Theory of Tikhonov Regularization for Fredholm Equations of the First Kind*, Pitman Publishing Ltd., London, 1984.
11. K. Miller, "Least-squares methods for ill-posed problems with a prescribed bound," SIAM J. Math. Anal., Vol. 1, No. 1, pp. 52-74, 1970.
12. A. Albert, *Regression and the Moore-Penrose Pseudoinverse*, Academic Press, New York and London, 1972.
13. G. Wahba, "Practical approximate solutions to linear operator equations when the data are noisy," *SIAM J. Numer. Anal.*, Vol. 14, No. 4, pp. 651-667, 1977.
14. P. C. Hansen, "Analysis of discrete ill-posed problems by means of the L-curve," *Argonne National Laboratory Preprint MCS-P157-0690*, July 1990.
15. G. H. Golub and C. Reinsch, "Singular value decomposition and least squares solutions," *Numer. Math.*, Vol. 14, pp. 403-420, 1970.
16. P. C. Hansen, "Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numerical rank," *SIAM J. Sci. Stat. Comput.*, Vol. 11, No. 3, pp. 503-518, May 1990.
17. P. J. Davis, *Circulant Matrices*, John Wiley & Sons, New York, 1979.
18. C.L.Matson, "Fourier spectrum extrapolation and enhancement using support-constraints," *Proc. SPIE*, 1767 (1992).
19. J. M. Speiser and H. J. Whitehouse, "Multidimensional hybrid signal processing architectures", *Proceedings of the SPIE Advanced Institute on Transformations in Optical Signal Processing*, SPIE Vol. 373, 1981.
20. G. Ammar and P. Gader, "New decompositions of the inverse of a Toeplitz matrix," in: A.A. Kaaschock, J. H. van Schuppen and A. C. M. Ran (Eds.), *Signal Processing, Scattering and Operator Theory, and Numerical Methods*, Birkhauser, Boston, 1990.
21. R. C. Gonzalez and P. Wintz, *Digital Image Processing* (Second Edition), Addison-Wesley, Reading, Massachusetts, 1987.
22. G. H. Golub and C. F. Van Loan, *Matrix Computations* (Second Edition), The Johns Hopkins University Press, Baltimore, 1989.
23. T. Kailath, S.Y. Kung and M. Morf, "Displacement ranks of matrices and linear equations," *J. Math. Anal. and Appl.*, Vol. 68, pp. 395-407, 1979.
24. J. M. Speiser, H. J. Whitehouse and J. C. Allen, "Fast matrix-vector multiplication using displacement rank approximation via an SVD", *Proceedings Second International Workshop on SVD and Signal Processing*, University of Rhode Island, 25-27 June 1990.

APPENDIX A - TRADEOFFS BETWEEN ACCURACY AND SPEED

In this appendix we examine the RAPID reconstruction scheme in the continuous domain. An expression for the error displays a compromise between accuracy and speed. Let A denote the convolution operator on $L^2(\Re^2)$:

$$Af(x) = (a \star f)(x) = \int_{\Re^2} a(x-y) f(y)\, dm(y).$$

Let $\hat{}$ denote the two-dimensional Fourier transform and let $M_{\hat{a}}$ denote the multiplication operator: $M_{\hat{a}}\varphi = \hat{a}\phi$. As is well-known (see reference A1 below), the Fourier transform diagonalizes the convolution operator: $\hat{}^{-1} A \hat{} = M_{\hat{a}}$, with $\|A\| = \|\bar{a}\|_\infty$. In this formalism, the regularization operator satisfies:

$$\hat{}^{-1} A_\beta^+ \hat{} = M_{(|\hat{a}|^2 + \beta)^{-1} \bar{\hat{a}}}.$$

The RAPID operator is obtained by determining the regularization pseudoinverse and then restricting it to a region near the origin. Introduce the associated convolution operator

$$A_b f(x) = \int_{\Re^2} (w_b a)(x-y) f(y)\, dm(y),$$

where w_b denotes the boxcar window: $w_b(x) = 1$ if $x \in [-b, b] \times [-b, b]$, zero otherwise. The regularization operator for A_β is

$$\hat{}^{-1} A_{b,\beta}^+ \hat{} = M_{w_b \star (|\hat{a}|^2 + \beta)^{-1} \bar{\hat{a}}}.$$

Then the error between the regularization operator and the RAPID operator is

$$\left\| A_\beta^+ - A_{b,\beta}^+ \right\| = \left\| \frac{\bar{\hat{a}}}{|\hat{a}|^2 + \beta} - w_b \star \frac{\bar{\hat{a}}}{|\hat{a}|^2 + \beta} \right\|_\infty.$$

As b tends to infinity, the error tends to zero. However, a large value for b corresponds to an increased computational load. Thus, a trade off must be made between speed and accuracy. In addition, the error bound also indicates that a different window (such as a square Hamming) might improve the accuracy of the regularization scheme.

A1. N. Young, *An Introduction to Hilbert Space*, Cambridge University Press, 1988.